

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	2	(parallel\$5 with hash\$3 with (join\$3 or merg\$3 or combin\$5)).ab. and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/16 20:33
S2	78	(hash\$3 with (join\$3 or merg\$3 or combin\$5)).ab. and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:01
S3	46	equijoin and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 07:51
S4	0	(online with aggregat\$3 with adapt\$3 with (quer\$5 or search\$3)).ab. and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 07:58
S5	1	(online with aggregat\$3 with (quer\$5 or search\$3)).ab. and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 08:00
S6	0	(online adj2 aggregat\$3 adj2 (quer\$5 or search\$3)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 08:09
S7	0	(parallel\$3 with (hash adj ripple adj join)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 08:10
S8	0	(hash adj ripple adj join) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:02
S9	6	(hash with join) same (partition\$3 with tuple) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/05/18 08:52
S10	30	(hash\$3 with (join\$3 or merg\$3 or combin\$5)).ab. and (@ad<"20010426") and partition\$3	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/11/14 00:42
S11	38	((((equivalent or hash\$3) with (join\$3 or merg\$3 or combin\$5)) with (processors or processers)) and partition\$3 and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2003/11/14 14:22
S12	10	((((equivalent or hash\$3) with (join\$3 or merg\$3 or combin\$5)) with (distribut\$3 or redistribut\$3) with (partition\$3 or divid\$3)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:23
S13	155	(hash with join\$3) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:06

**BEST AVAILABLE COPY**

S14	90	(hash with join\$3) and (aggregat\$5 or (intermed\$9 with result)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:14
S15	45	(hash with join\$3) and (((online or (on adj line)) adj aggregat\$5) or (intermed\$9 with result)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 16:20
S16	0	(hash with join\$3) and (split adj vector) and (((online or (on adj line)) adj aggregat\$5) or (intermed\$9 with result)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:19
S17	0	(hash with join\$3) and (split adj vector) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:19
S18	15	(hash with join\$3) and vector and (((online or (on adj line)) adj aggregat\$5) or (intermed\$9 with result)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 13:23
S19	41	(hash with join\$3) and vector and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/22 14:24
S20	0	(hash with join\$3) and (display\$3 with (((online or (on adj line)) adj aggregat\$5) or (intermed\$9 with result))) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/29 07:49
S21	1	"5884320".pn.	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/29 07:46
S22	35	((join or joining) with simultan\$9 with (distribut\$3 or redistribut\$3)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/29 07:55
S23	503	((join or joining) with (simultan\$9 or "as" or "while") with (distribut\$3 or redistribut\$3)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/29 08:00
S24	81	((join or joining) with (simultan\$9 or "as" or "while")) same (distribut\$3 or redistribut\$3)) and (parallel\$3 with process\$3) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2004/04/29 08:03
S25	1121	(hash\$3 with (join\$3 or merg\$3 or combin\$5)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:28
S26	26	(hash\$3 with (join\$3 or merg\$3 or combin\$5)) and ((join or joining or merg\$3 or combin\$3) with (simultan\$9 or "as" or "while") with (distribut\$3 or redistribut\$3)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:26

S27	30	((hash\$3 with (join\$3 or merg\$3 or combin\$5)) same (aggregat\$5 or (intermed\$9 with result))) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:27
S28	11	((hash\$3 with (join\$3 or merg\$3 or combin\$5)) same (((online or (on adj line)) adj aggregat\$5) or (intermed\$9 with result))) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:27
S29	42	(hash\$3 with (join\$3 or merg\$3 or combin\$5)) and (split\$5 with (vector or table or matrix)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:29
S30	2	(hash\$3 with (join\$3 or merg\$3 or combin\$5)) and ((distribut\$3 or redistribut\$3) with split\$5 with (vector or table or matrix)) and (@ad<"20010426")	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/14 20:29
S31	2	"6804678".pn. "6745198".pn.	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/17 14:20
S32	16	"4930072".pn. "5551031".pn. "5557791".pn. "5745896".pn. "5832475".pn. "6032144".pn. "6061676".pn. "6081801".pn. "6112198".pn. "6205441".pn. "6226639".pn. "6415297".pn. "6484159".pn. "6493701".pn. "6618719".pn. "6625593".pn. "22020103793"	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/17 18:06
S33	7	("4930072".pn. "5551031".pn. "5557791".pn. "5745896".pn. "5832475".pn. "6032144".pn. "6061676".pn. "6081801".pn. "6112198".pn. "6205441".pn. "6226639".pn. "6415297".pn. "6484159".pn. "6493701".pn. "6618719".pn. "6625593".pn. "22020103793") and (hash\$3 with join\$3) and (distribut\$3 or redistribut\$3) and partition\$3	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/17 16:18
S34	9	("4930072".pn. "5551031".pn. "5557791".pn. "5745896".pn. "5832475".pn. "6032144".pn. "6061676".pn. "6081801".pn. "6112198".pn. "6205441".pn. "6226639".pn. "6415297".pn. "6484159".pn. "6493701".pn. "6618719".pn. "6625593".pn. "22020103793") and intermediat\$3	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/17 16:07

S35	1	("4930072".pn. "5551031".pn. "5557791".pn. "5745896".pn. "5832475".pn. "6032144".pn. "6061676".pn. "6081801".pn. "6112198".pn. "6205441".pn. "6226639".pn. "6415297".pn. "6484159".pn. "6493701".pn. "6618719".pn. "6625593".pn. "22020103793") and (hash\$3 with join\$3) and ((dynamic\$5 with distribut\$3) or redistribut\$3) and partition\$3	US-PGPUB; USPAT; IBM_TDB	OR	ON	2005/02/17 16:18
-----	---	---	--------------------------------	----	----	------------------

File 347: JAPIO Nov 1976-2003/Dec(Updated 040402)

(c) 2004 JPO & JAPIO

File 350: Derwent WPIX 1963-2004/UD,UM &UP=200426

(c) 2004 Thomson Derwent

Set	Items	Description
S1	13903	(JOIN??? OR EQUIJOIN???) (7N) (ROW? ? OR TUPLE? ? OR COLUMN? ? OR FIELD? ? OR CELL? ? OR TABLE? ?)
S2	1295	S1(7N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARAL- LEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIME OR AS OR WHILE)
S3	478490	DATABASE? ? OR DATA()BASE? ? OR TABLE? ?
S4	204	S2 AND S3
S5	24	S4 AND IC=G06F
S6	363073	ROW? ? OR TUPLE? ? OR COLUMN? ?
S7	18246	S6(5N) (REDISTRIBUT? OR DISTRIBUT? OR SHUFFL? OR RESHUFFL? - OR SHIFT??? OR RESHIFT??? OR MOV??? OR MIX??? OR REORDER??? OR SORT??? OR RESORT??? OR REARRANG? OR REORGANI? OR REGROUP??? OR RE() (ARRANG? OR ORGANI? OR GROUP??? OR ORDER?))
S8	2645	S1(25N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARAL- LEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC - OR SAME()TIME OR AS OR WHILE)
S9	433	S1 AND S8 AND S3
S10	6	S9 AND S7
S11	4	S10 NOT S5
S12	17887	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (20N) (S- IMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNC- HRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIM- E)
S13	23	S12 AND S3 AND IC=G06F
S14	17	S13 NOT (S5 OR S11)

---

5/5/1 (Item 1 from file: 347)  
DIALOG(R)File 347:JAPIO  
(c) 2004 JPO & JAPIO. All rts. reserv.

02264412 \*\*Image available\*\*  
JOINT BODY RESEARCH SUPPORT SYSTEM

PUB. NO.: 2002-132873 [JP 2002132873 A]  
PUBLISHED: May 10, 2002 (20020510)  
INVENTOR(s): BEPPU FUSAO  
APPLICANT(s): AINEKKUSU KK  
APPL. NO.: 2000-358501 [JP 2000358501]  
FILED: October 20, 2000 (20001020)  
INTL CLASS: G06F-017/60 ; G06F-019/00

#### ABSTRACT

PROBLEM TO BE SOLVED: To provide a joint body research support system which can improve the reliability of a result body to be delivered to an electric power company, reduce operation for entry into a joint facility research table by a field operator making a research into joint bodies, and shorten the research time.

SOLUTION: An area power distribution block database is generated by inputting joint body information based upon the contents of a contract with a requester for the installation of joint bodies by area power distribution blocks as objects of research and the joint body facility research table is generated which is used to make a previous research of joint body articles, inspect differences from an actual state after the field research is completed, totalize information of contents different from joint body information based upon the contract with the mentioned requester, and generate the result body by inputting information generated by judging a joint cable of a joint body found by the field research by a communication cable external-diameter measurement system and entering all research results.

COPYRIGHT: (C)2002,JPO

5/5/2 (Item 2 from file: 347)  
DIALOG(R)File 347:JAPIO  
(c) 2004 JPO & JAPIO. All rts. reserv.

04062383 \*\*Image available\*\*  
DATABASE RETRIEVAL SYSTEM

PUB. NO.: 05-054083 [JP 5054083 A]  
PUBLISHED: March 05, 1993 (19930305)  
INVENTOR(s): ICHIYAMA TOSHIHARU  
APPLICANT(s): NEC CORP [000423] (A Japanese Company or Corporation), JP  
(Japan)  
APPL. NO.: 03-235452 [JP 91235452]  
FILED: August 21, 1991 (19910821)  
INTL CLASS: [5] G06F-015/40 ; G06F-012/00  
JAPIO CLASS: 45.4 (INFORMATION PROCESSING -- Computer Applications); 45.2  
(INFORMATION PROCESSING -- Memory Units)  
----- JOURNAL ----- Section: P, Section No. 1571, Vol. 17, No. 366, Pg. 14, July  
09, 1993 (19930709)

#### ABSTRACT

PURPOSE: To automatically extract the joint information between tables without any leaking even in the database composed of many tables by retrieving the database only by describing the conditions concerning the column of respective tables even when a user does not have the knowledge concerning the structure of the database which is a retrieval object and the detailed knowledge concerning the joint information between the tables.

CONSTITUTION: This system has a joint information holding means 4 to hold

the joint information between tables of a database, a joint conditions imparting means 1 to obtain the joint information from the joint information holding means 4 and output a retrieval expression to add the joint information with the table retrieval expression as an input and a database control means 2 to perform the retrieval of a database storage means 3 while the retrieval expression is received, and has a joint information extraction means 5 to extract the joint information between the tables of the database storage means 3 from the database control means 2 and send the extracted joint information to the joint information holding means 4.

5/5/3 (Item 3 from file: 347)  
... 347:JAPIO  
... JAPIO. All rts. reserv.

141861 \*\*Image available\*\*  
OUTER JOINING OPERATION SYSTEM OF RELATIONAL DATA BASE

PUB. NO.: 59-125461 [JP 59125461 A]  
PUBLISHED: July 19, 1984 (19840719)  
INVENTOR(s): TEZUKA MASAYOSHI  
ADACHI SUSUMU  
NAKADA TERUO  
YAMANE YASUO  
MAKINOCHI AKIFUMI  
APPLICANT(s): FUJITSU LTD [000522] (A Japanese Company or Corporation), JP  
(Japan)  
APPL. NO.: 57-234066 [JP 82234066]  
FILED: December 30, 1982 (19821230)  
INTL CLASS: [3] G06F-013/00 ; G06F-007/22  
JAPIO CLASS: 45.2 (INFORMATION PROCESSING -- Memory Units); 45.1  
(INFORMATION PROCESSING -- Arithmetic Sequence Units)  
JOURNAL: Section: P, Section No. 315, Vol. 08, No. 254, Pg. 87,  
November 21, 1984 (19841121)

#### ABSTRACT

PURPOSE: To realize a high-speed processing system by scanning simultaneously relation to be joined together, comparing to decide whether a join field value is larger, smaller, or is equal, etc., to perform outer joining, and to produce the relation of result.

DEFINITION: A relational data base processing mechanism 1 accepts an inquiry from a user and passes it to a relational operation control part 2. The relational operation control part 2 checks the adequacy of the inquiry. An execution control part 6 sorts and extracts relation operation to determine its execution procedure, which is passed to an interpretation execution part 3 and executed. Then, an outer joining execution part 10 performs outer joining processing. In outer join operation processing, the values of joint field are extracted and compared successively from the starting record of plural sort relation to generate outer joining result responding to the result decided to be larger, smaller or is equal, and the result is inserted into the relation of results.

---

5/5/4 (Item 1 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

015892836 \*\*Image available\*\*  
WPI Acc No: 2004-050671/200405  
XRPX Acc No: N04-040933

Query optimization method in relational database system, involves joining pair of tables having respective primary and foreign key columns using specified join condition, and generating derived data constraint rule

Parent Assignee: NCR CORP (NATC )  
Inventor: GHAZAL A S; SINCLAIR P L

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 6662175	B1	20031209	US 2001850765	A	20010508	200405 B

Priority Applications (No Type Date): US 2001850765 A 20010508

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 6662175	B1	15	G06F-017/30	

Abstract (Basic): US 6662175 B1

NOVELTY - A pair of **tables** having respective primary key (PK) and foreign key (FK) **columns** are **joined**, using PK is equal to FK as join condition. An initial running constraint (RC) having null range is created. A derived date constraint rule (DDCR) is generated, based on correlated value columns of join result. The RC is modified by merging range of RC with range of new constraint that is computed for each row in join result.

DETAILED DESCRIPTION - An INDEPENDENT CLAIM is also included for computer program for optimizing queries to **database**.

USE - For optimizing queries in relational **database** system such as teradata active data warehousing system and other **database** systems having massively parallel processing (MPP) architecture or symmetric multiprocessing (SMP) architecture.

ADVANTAGE - Automatically derives constraint rules for correlated variables, and reduces the size of intermediate result.

DESCRIPTION OF DRAWING(S) - The figure shows a flowchart of algorithm for deriving and applying derived date constraint rule.  
pp; 15 DwgNo 4/9

Title Terms: QUERY; OPTIMUM; METHOD; RELATED; **DATABASE** ; SYSTEM; JOIN;

PAIR; **TABLE** ; RESPECTIVE; PRIMARY; FOREIGN; KEY; COLUMN; SPECIFIED; JOIN ; CONDITION; GENERATE; DERIVATIVE; DATE; CONSTRAIN; RULE

Derwent Class: T01

International Patent Class (Main): **G06F-017/30**

File Segment: EPI

5/5/5 (Item 2 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

015834435 \*\*Image available\*\*

WPI Acc No: 2003-896639/200382

Related WPI Acc No: 2001-496729; 2003-660928; 2003-897748

XRPX Acc No: N03-715546

**Data tubes populate method for data management system, involves joining columns or domains of table with dimensions and other relations mapped into hypercube to populated hypercube**

Patent Assignee: DECODE GENETICS EHF (DECO-N)

Inventor: EGILSSON A S; GUDBJARTSSON H

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 20030023608	A1	20030130	US 99475436	A	19991230	200382 B
			US 2002216670	A	20020808	

Priority Applications (No Type Date): US 2002216670 A 20020808; US 99475436 A 19991230

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 20030023608	A1	27	G06F-007/00	CIP of application US 99475436 CIP of patent US 6434557

Abstract (Basic): US 20030023608 A1

NOVELTY - The method involves representing a calculated relation as a **table** supported by **columns** or domains and joining **columns** or domains of the **table** with dimensions and other relations mapped into a hypercube. The new relations are created from existing relations

and table-like representations of calculated relations.

DETAILED DESCRIPTION - An INDEPENDENT CLAIM is also included for data tube populate system.

USE - For data management system.

ADVANTAGE - Eliminates ambiguities observed in combined measurements used to populate a hypercube.

DESCRIPTION OF DRAWING(S) - The figure shows a block diagram of the data tube populate system.

data tube populate system (100)

database (101)

online analytic processing server (102)

online analytic processing client (103)

network (104)

pp; 27 DwgNo 1/14

Title Terms: DATA; TUBE; METHOD; DATA; MANAGEMENT; SYSTEM; JOIN; COLUMN;  
DOMAIN; TABLE; DIMENSION; RELATED; MAP; POPULATION

Derwent Class: T01

International Patent Class (Main): G06F-007/00

Derwent: EPI

5/5/6 (Item 3 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

015397411 \*\*Image available\*\*

WPI Acc No: 2003-459551/200344

XRAM Acc No: C03-122363

XRPX Acc No: N03-365537

Safety busbar system, especially for a tablet production machine  
comprises system elements which are jointly capable of simultaneous  
implementation of safety, control and measurement functions

Patent Assignee: KORSCH AG (KORS-N)

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
DE 20304950	U1	20030528	DE 2003U2004950	U	20030321	200344 B

Priority Applications (No Type Date): DE 2003U2004950 U 20030321

Patent Details:

Patent No	Kind	Lan	Pg	Main IPC	Filing Notes
DE 20304950	U1	14	G06F-013/38		

Abstract (Basic): DE 20304950 U1

NOVELTY - Safety busbar system (10), especially for a tablet  
production machine (12) comprises at least one module (14) compatible  
with the busbar system and/or at least one busbar controller (16) and  
at least one busbar conductor (18). The system ensures essentially  
simultaneous implementation of at least one safety function, at least  
one control function and/or at least one measurement function.

USE - Used as a safety system for a tablet production machine.

ADVANTAGE - The amount of wiring and the general system complexity  
are reduced together with fault hazards.

DESCRIPTION OF DRAWING(S) - The drawing shows a schematic view of  
the proposed safety busbar system.

~~Tablet production machine (12)~~

Module (14)

Busbar controller (16)

Busbar conductor (18)

Actuator (20)

Sensor (22)

Emergency disconnect switch (24)

pp; 14 DwgNo 1/1

Title Terms: SAFETY; SYSTEM; TABLET; PRODUCE; MACHINE; COMPRISE; SYSTEM;  
ELEMENT; JOINT; CAPABLE; SIMULTANEOUS; IMPLEMENT; SAFETY; CONTROL;  
MEASURE; FUNCTION

Derwent Class: B07; X12; X25

International Patent Class (Main): G06F-013/38

File Segment: CPI; EPI

5/5/7 (Item 4 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

015378530 \*\*Image available\*\*  
WPI Acc No: 2003-439468/200341  
XRPX Acc No: N03-350664

Document mining method involves removing instance of dirty text within document to produce clean document and then performing data mining operation

Patent Assignee: CASTELLANOS M (CAST-I); STINGER J R (STIN-I)

Inventor: CASTELLANOS M; STINGER J R

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 20030046263	A1	20030306	US 2001944919	A	20010831	200341 B

Priority Applications (No Type Date): US 2001944919 A 20010831

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 20030046263	A1	13	G06F-007/00	

Abstract (Basic): US 20030046263 A1

NOVELTY - An instance of dirty text within the document, is removed to produce a clean document. Data mining operation is then performed on the clean document.

DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following:

(1) computer system; and  
(2) computer usable recorded medium storing document mining program.

USE - For mining document containing dirty text such as typographical errors, misspellings, joined words, ad hoc abbreviations, bad grammar, cryptic tables, programming code, core dumps, missing or ambiguous punctuation and haphazard capitalization.

ADVANTAGE - Allows the user to leverage existing domain knowledge and enables easy customization of document containing dirty text according to the domain and task requirements.

DESCRIPTION OF DRAWING(S) - The figure shows the block diagram of the document mining system.

pp; 13 DwgNo 2/5

Title Terms: DOCUMENT; MINE; METHOD; REMOVE; INSTANCE; DIRT; TEXT; DOCUMENT ; PRODUCE; CLEAN; DOCUMENT; PERFORMANCE; DATA; MINE; OPERATE

Derwent Class: T01

International Patent Class (Main): G06F-007/00

File Segment: EPI

5/5/8 (Item 5 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

~~2002-626900~~ \*\*Image available\*\*  
WPI Acc No: 2002-626900/200267  
Related WPI Acc No: 2002-425217  
XRPX Acc No: N02-495811

Replication system for database management in firm, matches data items in several master tables using preset data in the tables and generates replica table by replicating the matched data items

Patent Assignee: YOKOUCHI H (YOKO-I)

Inventor: YOKOUCHI H

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 20030046263	A1	20020711	US 2001809257	A	20010316	200267 B

Priority Applications (No Type Date): JP 2000294551 A 20000927

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes  
US 20020091716 A1 14 G06F-007/00 Cont of application US 2001809257

Abstract (Basic): US 20020091716 A1

NOVELTY - The data items in several master tables (108,109) are matched using preset data in the master tables as a joining key. A replica table (121) is generated by replicating the matched data items of the master tables.

DETAILED DESCRIPTION - An INDEPENDENT CLAIM is included for replication program.

USE - Replication system for database including salary details in a firm.

ADVANTAGE - As the data items of several master tables are replicated and included in a single replica table, job execution time of the replica system, is reduced, thus efficiency is improved.

DESCRIPTION OF DRAWING(S) - The figure shows the data flow in the replication system.

Master tables (108,109)

Replica table (121)

pp: 14 DwgNo 5/7

Title Terms: REPLICAS; SYSTEM; DATABASE; MANAGEMENT; FIRM; MATCH; DATA; ITEM; MASTER; TABLE; PRESET; DATA; TABLE; GENERATE; REPLICAS; TABLE; REPLICAS; MATCH; DATA; ITEM

Derwent Class: T01

International Patent Class (Main): G06F-007/00

File Segment: EPI

5/5/9 (Item 6 from file: 350)

MAILLOG(R)File 350:Derwent WPIX

1/1/2004 Thomson Derwent. All rts. reserv.

014604513 \*\*Image available\*\*

WPI Acc No: 2002-425217/200245

Related WPI Acc No: 2002-626900

XRPX Acc No: N02-334351

Data replication system for employee management system, matches data items of several master tables and replicates matched data items to generate replica table using particular data in master table

Patent Assignee: HITACHI LTD (HITA); YOKOUCHI H (YOKO-I)

Inventor: YOKOUCHI H

Number of Countries: 002 Number of Patents: 002

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 20020038315	A1	20020328	US 2001809257	A	20010316	200245 B
JP 2002108681	A	20020412	JP 2000294551	A	20000927	200245

Priority Applications (No Type Date): JP 2000294551 A 20000927

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes

US 20020038315 A1 14 G06F-017/30

JP 2002108681 A 12 G06F-012/00

Abstract (Basic): US 20020038315 A1

NOVELTY - A matching unit matches data items of a number of master tables (108,109) using particular data in the master tables as a joining key to join each other. A replicating unit replicates the matched data items based on the joining key to generate one replica table (121).

DETAILED DESCRIPTION - An INDEPENDENT CLAIM is also included for data replication program.

USE - For replicating data of several master tables of database in employee management system.

ADVANTAGE - The data of number of master tables is replicated in one replica table, and hence job execution time is minimized.

DESCRIPTION OF DRAWING(S) - The figure shows the block diagram of data replica system.

Master tables (108,109)

Replica table (121)

pp; 14 DwgNo 1/7

Title Terms: DATA; REPLICA; SYSTEM; EMPLOY; MANAGEMENT; SYSTEM; MATCH; DATA ; ITEM; MASTER; TABLE ; REPLICA; MATCH; DATA; ITEM; GENERATE; REPLICA; TABLE ; DATA; MASTER; TABLE

Derwent Class: T01

International Patent Class (Main): G06F-012/00 ; G06F-017/30

File Segment: EPI

5/5/10 (Item 7 from file: 350)

5 5/5/10: Derwent WPIX

4 Thomson Derwent. All rts. reserv.

01445412 \*\*Image available\*\*

WPI Acc No: 2002-276115/200232

XRPX Acc No: N02-215577

Elbow rest has joint that is fitted to base, fixing tool that fixes joint to base, support that protrudes in front of base when joint is fixed to base, and elbow accepting unit that is attached to support

Patent Assignee: HOGURA Y (HOGU-I)

Number of Countries: 001 Number of Patents: 001

Patent Family:

Parent No	Kind	Date	Applicat No	Kind	Date	Week
JP 2002051862	A	20020219	JP 2000245062	A	20000811	200232 B

Priority Applications (No Type Date): JP 2000245062 A 20000811

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes

JP 2002051862 A 8 A47B-096/18

Abstract (Basic): JP 2002051862 A

NOVELTY - A fixing tool (3) fixes a joint (2) that is fitted to a base (1) such as a desk or table. A support articulated by the joint protrudes in front of the base when the joint is fixed to the base. An elbow accepting unit (5) is attached to the support.

USE - For e.g. desk, table.

ADVANTAGE - Simplifies attachment to or detachment from base. Enables height to be adjusted. Enables work bench to be reduced in size.

DESCRIPTION OF DRAWING(S) - The figure is the perspective diagram of the elbow rest.

Base (1)

Joint (2)

Fixing tool (3)

Elbow accepting unit (5)

pp; 8 DwgNo 1/9

Title Terms: ELBOW; REST; JOINT; FIT; BASE; FIX; TOOL; FIX; JOINT; BASE; SUPPORT; PROTRUDE; FRONT; BASE; JOINT; FIX; BASE; ELBOW; ACCEPT; UNIT; ATTACH; SUPPORT

Derwent Class: P25; T01

International Patent Class (Main): A47B-096/18

International Patent Class (Additional): A47B-011/00; A47B-017/02;

A47B-017/03; A47B-037/00; G06F-001/16 ; G06F-003/02

File Segment: EPI; EngPI

5/5/11 (Item 8 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

014094033

WPI Acc No: 2001-578247/200165

XRPX Acc No: N01-430163

Relational database systems, using equijoin operations giving a merge

join process that creates sorted set of outer table rows that satisfy selection criteria and enable parallel execution

Patent Assignee: UNISYS CORP (BURS )

Inventor: LIU L H

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 6185557	B1	20010206	US 98135312	A	19980731	200165 B

Priority Applications (No Type Date): US 98135312 A 19980731

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 6185557	B1	13	G06F-017/00	

Abstract (Basic): US 6185557 B1

NOVELTY - When performing **equijoin** operations on two **tables**, as long as outer **table join column** value is less than or equal to last key value, the same data page from inner **table** will be searched repeatedly. Therefore inner **table** index records will not be revisited until outer **table join column** is greater than last data page key value. Combination of next key and last key allows merge join process to determine that entire ranges of outer rows do not have matching inner rows.

DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following:

(1) A computer readable medium having computer executable modules for joining data rows from two **tables**.

(2) A merge join process for joining data rows from two **tables** which have a common data column.

USE - Relational **database** systems

ADVANTAGE - Designed to minimize processor time and file input and output when performing **equijoin** operations on two **tables**. Reduces or eliminates searches of the inner index records and the data pages. Reduces cache thrashing on the inner index records and so the required index record is likely to be in the cache when needed. Minimizing the traversal of the index records and data pages on mass storage minimizes the number of operations performed, and therefore provides a more efficient search process. Also, because the merge join process is structurally suited for execution on the multi-processor computers, the speed of the **database** queries can be increased through parallel processing.

DESCRIPTION OF DRAWING(S) - Merge join process logic flow diagram.

pp; 13 DwgNo 0/4

Title Terms: RELATED; **DATABASE**; SYSTEM; OPERATE; MERGE; JOIN; PROCESS; SORT; SET; OUTER; **TABLE**; ROW; SATISFY; SELECT; CRITERIA; ENABLE; PARALLEL; EXECUTE

Derwent Class: T01

International Patent Class (Main): G06F-017/00

File Segment: EPI

5/5/12 (Item 9 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

-----013690202-----\*\*Image available\*\*-----

WPI Acc No: 2001-174426/200118

XRPX Acc No: N01-126308

Database outputs required search data by joining data retrieved from tables specified by divided main keys

Patent Assignee: SONY CORP (SONY )

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
JP 2001005815	A	20010112	JP 99171007	A	19990617	200118 B

Priority Applications (No Type Date): JP 99171007 A 19990617

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes  
JP 2001005815 A 7 G06F-017/30

Abstract (Basic): JP 2001005815 A

NOVELTY - The main key corresponding to each data is divided according to components of the data such that each divided main key specifies the **table** (120A-120C). In order to search data, the main key input by operator is divided and data in each **table** specified by the divided main key is **joined** and output as the required result.

USE - In **database** for searching data from **tables**.

ADVANTAGE - The search time can be shortened by using the divided main key instead of unitary main key for searching purpose. Due to the reduction of amount of data and shortening of search time, the workload of maintenance operation can also be reduced.

DESCRIPTION OF DRAWING(S) - The figure shows explanatory drawing of **database**.

Table (120A-120C)

Fig. 7 DwgNo 1/3

Forms: DATABASE ; OUTPUT; REQUIRE; SEARCH; DATA; JOIN; DATA;

RETRIEVAL; TABLE ; SPECIFIED; DIVIDE; MAIN; KEY

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/13 (Item 10 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

013575313 \*\*Image available\*\*

WPI Acc No: 2001-059520/200107

XRPX Acc No: N01-044405

Incremental refresh performing method for materialized view in database management systems, involves deleting rows, that are attained by combining specific rows with changed rows of selected table

Patent Assignee: ORACLE CORP (ORAC-N)

Inventor: DIAS K; WITKOWSKI A

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 6125360	A	20000926	US 98109115	A	19980702	200107 B

Priority Applications (No Type Date): US 98109115 A 19980702

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes

US 6125360 A 24 G06F-017/30

Abstract (Basic): US 6125360 A

NOVELTY - The base **table** of materialized view is established as selected **table**. The row (Tj) which is processed, is combined with both changed and unchanged rows of selected **table**. The rows that are attached due to combining of changed rows of selected **table** with Tj in materialized view is deleted, and the rows in which Tj combines with unchanged rows of selected **table** is left.

DETAILED DESCRIPTION - The base **table** of a materialized view is established as a selected **table**. If the selected **table** is the right **table** of outer join, the materialized view is updated to reflect deletions to selected **table** by processing each row (Tj), that combines with changed row of selected **table**. If Tj combines in materialized view with changed rows of selected **table**, then the rows containing Tj in materialized view is removed and replaced with a row in which selected columns are set to NULL. INDEPENDENT CLAIMS are also included for the following:

(a) **database** system;

(b) incremental refresh program

USE - For maintenance of materialized views that contain one-to-N lossless joins in **database** management system.

ADVANTAGE - Since the incremental refresh technique does not

require information about the order of updates to the base tables , the overhead associated with maintaining sequencing information is avoided. Since the technique is idempotent in performing an incremental refresh N-times on the same data, after a system crash, the incremental refresh operation is restarted from the beginning without taking into account how far the operation had progressed prior to the crash.

DESCRIPTION OF DRAWING(S) - The figure shows the flowchart that illustrates the incremental refresh operation.

pp: 24 DwgNo 5/7

Title Terms: INCREMENT; REFRESH; PERFORMANCE; METHOD; VIEW; DATABASE ;  
MANAGEMENT; SYSTEM; DELETE; ROW; ATTAIN; COMBINATION; SPECIFIC; ROW;  
CHANGE; ROW; SELECT; TABLE

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/14 (Item 11 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

013526696 \*\*Image available\*\*

WPI Acc No: 2001-010902/200102

KRPX Acc No: N01-008368

Computer aided designing/computer aided machine conversion system e.g.  
for PCB fabrication, converts CAD system data to source table , inorder  
to link data edition number with comprehensive edition number

Patent Assignee: FUJITSU LTD (FUIT )

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
JP 2000263354	A	20000926	JP 9974476	A	19990318	200102 B

Priority Applications (No Type Date): JP 9974476 A 19990318

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
JP 2000263354	A	17	B23P-021/00	

Abstract (Basic): JP 2000263354 A

NOVELTY - Data of each CAD system contains intrinsic format and data edition number with format common to a CAD system. A converter converts CAD system data to source table comprising information regarding board, mounting, node, inorder to link data editing number with comprehensive edition number which are managed by a management unit.

DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following:

(a) CAD/CAM conversion procedure;

(b) recording medium with data conversion program

USE - E.g. for printed circuit board (PCB) fabrication.

ADVANTAGE - Automatic conversion is performed at mounting angle joined on output destination line as table maintains mounting angle information at both sides rotations for each component. Since data format of different CAD system is converted into an unification data format, each CAD data is managed with identical database .

----- DESCRIPTION OF DRAWING(S) ----- The figure shows the diagram illustrating components of CAD/CAM conversion system.

- pp: 17 DwgNo 1/17

Title Terms: COMPUTER; AID; DESIGN; COMPUTER; AID; MACHINE; CONVERT; SYSTEM ; PCB; FABRICATE; CONVERT; CAD; SYSTEM; DATA; SOURCE; TABLE ; LINK; DATA ; EDIT; NUMBER; COMPREHENSIVE; EDIT; NUMBER

Derwent Class: P56; T01; V04

International Patent Class (Main): B23P-021/00

International Patent Class (Additional): G06F-017/50 ; H05K-003/00

File Segment: EPI; EngPI

5/5/15 (Item 12 from file: 350)

DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

013493361 \*\*Image available\*\*

WPI Acc No: 2000-655304/200063

RRPX Acc No: N00-485721

Access management method for shared resource in multiprocessing system,  
involves choosing subject processor to operate in DUAL mode or SOLO mode  
in response to information exchange attempt being success or failure

Patent Assignee: INT BUSINESS MACHINES CORP (IBMC )

Inventor: FREITAS R F; JADAV D; KENCHAMMANA-HOSEKOTE D; MENON J M; STRONG H  
R

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 6105099	A	20000815	US 98203102	A	19981130	200063 B

Priority Applications (No Type Date): US 98203102 A 19981130

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 6105099	A		22 G06F-012/00	

Abstract (Basic): US 6105099 A

NOVELTY - Each of the processors (102,104) sends self-state message repeatedly to other processor. The subject processor chooses operating mode in response to preset power-up conditions. Any information required to synchronize processor's lock tables is exchanged suitably during join operation. If attempted information exchange succeeds, subject processor is chosen to operate in DUAL mode, else in SOLO mode.

DETAILED DESCRIPTION - The subject processor establishes message content as birth cry message when lock tables (108,110) of subject processor requires information from other processor's lock table. The processor is operated in DUAL mode or SOLO mode only after completion of join operation. When the processor in DUAL mode, the message content is established by processor as heart beat message. The processor combines with other processor to establish single processor locks on sub-parts of shared resources (106) to satisfy host requests to access the shared resource comprising digital data storage. When failure is indicated from other processor, the subject processor is operated in SOLO mode during which message content is established as death knell message. The processor attempts to acquire sub-part of shared resource and to store a predetermined code at each sub-part. The access to shared resource is enabled without consulting the other processor only by requirement of sub-parts of shared resource. INDEPENDENT CLAIMS are included for the following:

- (a) access management program;
- (b) digital data processing system

USE - For managing access to shared resources in multiprocessing systems used in data processing applications such as automated teller network, airline reservation systems, stock brokerage, etc.

ADVANTAGE - Makes better use of high through-put shared resource by efficiently sharing the resources among hierarchically superior hosts.

DESCRIPTION OF DRAWING(S) - The figure shows the block diagram of computer system with multiple processors.

-----  
Processors (102,104)

Shared resources (106)

Lock tables (108,110)

pp: 22 DwgNo 1A/8

Title Terms: ACCESS; MANAGEMENT; METHOD; SHARE; RESOURCE; MULTIPROCESSOR;  
SYSTEM; CHOICE; SUBJECT; PROCESSOR; OPERATE; DUAL; MODE; SOLO; MODE;  
RESPOND; INFORMATION; EXCHANGE; ATTEMPT; SUCCESS; FAIL

Derwent Class: T01

International Patent Class (Main): G06F-012/00

International Patent Class (Additional): G06F-012/14

File Segment: EPI

5/5/16 (Item 13 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

012672181 \*\*Image available\*\*  
WPI Acc No: 1999-478288/199940  
XRPX Acc No: N99-356022

Duplicate tuples elimination method in database management system

Patent Assignee: SYBASE INC (SYBA-N)

Inventor: HILLEGAS R

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 96757367	A	19990810	US 96757367	A	19961127	199940 B

Priority Applications (No Type Date): US 96757367 A 19961127

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 5937401	A	15	G06F-017/30	

Abstract (Basic): US 5937401 A

NOVELTY - The query for generating the tuple stream satisfying the selection criteria is executed, by scanning the selected database tables (250) according to the determined join order. The inner most table is scanned, thereby executing the filter which filters the duplicate tuples from the tuple stream.

DETAILED DESCRIPTION - A received query specifies the selection criteria for the information of interest from the database system. The determined join order indicates the innermost and outermost tables of the selected join so as to guarantee that the tuples will stream in order during scanning of the query. A filter is initialized at the outermost table for key columns to pass the initial tuple encountered from which an initial key is constructed. On execution the tuples having keys already encountered in the tuple stream are discarded by the filter attached to the innermost table.

USE - For eliminating duplicate tuples in a generated tuple stream in a database management system.

ADVANTAGE - The duplicate tuples are eliminated from the tuple stream without the need for performing an expensive sort operation by the described method.

DESCRIPTION OF DRAWING(S) - The figure is a block diagram of a client-server system with the duplicate tuples elimination method.

Database tables (250)

pp: 15 DwgNo 2/2

Title Terms: DUPLICATE; ELIMINATE; METHOD; DATABASE ; MANAGEMENT; SYSTEM

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/17 (Item 14 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

012325563 \*\*Image available\*\*  
WPI Acc No: 1999-131670/199911  
XRPX Acc No: N99-095967

Data processing system for accessing database - has catalogue analysis module which selects table within database, to create ID key for each set of columns in table identified as unique set

Patent Assignee: ACTUATE SOFTWARE INC (ACTU-N)

Inventor: YOUNG C

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 95426788	A	19990126	US 95426788	A	19950421	199911 B

Priority Applications (No Type Date): US 95426788 A 19950421

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes  
US 5864856 A 18 G06F-017/30

Abstract (Basic): US 5864856 A

NOVELTY - A dictionary builder (50) includes catalogue analysis module (52) that selects table within the database, and creates ID key for each set of one or more columns in the table identified as a unique set. A join key referencing selected table is created in any other table in the database having a set of columns with column names matching with that of ID key. DETAILED DESCRIPTION - A user interface (114) is coupled to a CPU (110). A dictionary buffer which is also coupled to CPU scans the database (130) to create dictionary containing attributes that define relationship between tables in the database and couples dictionary to the database. The table includes columns having column names. INDEPENDENT CLAIMS are also included for the following: information accessing method from database; dictionary builder.

USE - For accessing information from relational databases.

ADVANTAGE - Avoids need for extensive programming by system administrator to access database interactively. DESCRIPTION OF DRAWING(S) - The figure shows the block diagram of data processing system. (50) Dictionary builder; (52) Catalogue analysis module; (110) CPU; (114) User interface; (130) Database.

Dwg.2/11

Title Terms: DATA; PROCESS; SYSTEM; ACCESS; DATABASE; CATALOGUE; ANALYSE; MODULE; SELECT; TABLE; DATABASE; ID; KEY; SET; COLUMN; TABLE; IDENTIFY; UNIQUE; SET

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/18 (Item 15 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

012193669 \*\*Image available\*\*

WPI Acc No: 1998-610582/199851

XRPX Acc No: N98-474902

Joining method in database for one or more input tables comprising records stored in storage medium - using join index and minimizing number of input/output operations while maximizing use of small main memory through buffer allocation process based on join index entries

Patent Assignee: UNIV COLUMBIA NEW YORK (UYCO)

Inventor: LEI H; ROSS K A

Number of Countries: 021 Number of Patents: 002

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
WO 9850867	A1	19981112	WO 98US8339	A	19980424	199851 B
US 983215	A	19991109	US 97853108	A	19970508	199954

Priority Applications (No Type Date): US 97853108 A 19970508

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes

WO 9850867 A1 E 59 G06F-017/30

Designated States (National): CA JP

Designated States (Regional): AT BE CH CY DE DK ES FI FR GB GR IE IT LU

MC NL PT SE

US 983215 A G06F-017/30

Abstract (Basic): WO 9850867 A

The method involves allocating an array of partitions in a memory. Next a join index is read indicative of records to be joined. The join index entries are stored corresponding to each of one or more input tables of stored records and a corresponding partition identifier to temporary files associated with the allocated partitions. The index entries and the corresponding partition identifier in each of the

temporary files are read in turn and each temporary file is sorted. Portions of the tables are sequentially read only if the portion includes a record identified in the sorted temporary file.

The read records are written in accordance with an order of the corresponding partition identifiers to separate output files associated with each input table. The allocation step equally distributes the records of the input table among the array of partitions.

ADVANTAGE - Join technique performs join operation irrespective of join index order. Processes all input tables simultaneously. Can perform self-join operation without reading input table multiple times.

Dwg.3/9

Title Terms: JOIN; METHOD; DATABASE ; ONE; MORE; INPUT; TABLE ; COMPRISE; RECORD; STORAGE; STORAGE; MEDIUM; JOIN; INDEX; NUMBER; INPUT; OUTPUT; OPERATE; MAIN; MEMORY; THROUGH; BUFFER; ALLOCATE; PROCESS; BASED; JOIN; INDEX; ENTER

IPC Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/19 (Item 16 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

011915880 \*\*Image available\*\*

WPI Acc No: 1998-332790/199829

XRPX Acc No: N98-259808

Aid preparation method for recognising database statement with data structure in memory device of computer system - involves determining ratio of number of distinct rows satisfying join statement in detail table and master table for each directional link associated with detail and master table

Patent Assignee: ORACLE CORP (ORAC-N)

Inventor: TOW D S

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 5761654	A	19980602	US 96659158	A	19960605	199829 B

Priority Applications (No Type Date): US 96659158 A 19960605

Patent Details:

Patent No	Kind	IPC	Filing Notes
US 5761654	A	17 G06F-017/30	

Abstract (Basic): US 5761654 A

The method involves defining a set of nodes representing tables. A set of directional links are defined between pairs of nodes, each of which representing a master detail relationship between a detailed table and corresponding master table. A data structure referred to as a join free comprising representation of nodes and their directional links is defined in a memory device. A set of zeros or more selectivity factors are represented for each node.

Each selectivity factor indicates expected fraction of rows in table represented by node that satisfies one or more logical conditions set forth in the data access statement. The ratio of number of distinct rows satisfying join statement in detail table to number of distinct rows satisfying join statement in master table is determined for each directional link associated with detail table and master table. The probability of row in the detail table corresponds to that of rows in master table.

ADVANTAGE - Improves statement execution efficiency. Produces truly optimise statements.

Dwg.7/7

Title Terms: AID; PREPARATION; METHOD; RECOGNISE; DATABASE ; STATEMENT; DATA; STRUCTURE; MEMORY; DEVICE; COMPUTER; SYSTEM; DETERMINE; RATIO; NUMBER; DISTINCT; ROW; SATISFY; JOIN; STATEMENT; DETAIL; TABLE ; MASTER; TABLE ; DIRECTION; LINK; ASSOCIATE; DETAIL; MASTER; TABLE

Derwent Class: T01  
International Patent Class (Main): G06F-017/30  
File Segment: EPI

5/5/20 (Item 17 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

011988335 \*\*Image available\*\*  
WPI Acc No: 1998-322214/199828  
XRPX Acc No: N98-252044

Query optimisation method for relation database - involves finding JOIN conditions to form chain to form tables according to graph join theory and reordering in FROM clause

Patent Assignee: BULL HN INFORMATION SYSTEMS INC (HONE )

Inventor: GRAY J E

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 5758335	A	19980526	US 96722825	A	19960927	199828 B

Priority Applications (No Type Date): US 96722825 A 19960927

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 5758335	A	7	G06F-017/30	

Abstract (Basic): US 5758335 A

The query optimisation method involves examining the WHERE clause and finding all join conditions that are present. A chain of the join conditions is formed.

A list of tables is established where i) tables that are in the WHERE clause but not in the join conditions are listed first; ii) tables in the join chains formed according to graph theory are listed next in the same order as they appear in the join chain such that the tables from the longest chains are listed before the tables from shorter chains and all tables at a given distance from a root table or the chain occur together before the next level in the join chain. The tables are reordered in the FROM clause in the list order.

ADVANTAGE - Quickly finds access plan. Orders tables in FROM clause according to optimal join order in WHERE clause determined by graph theory.

Dwg.1/1

Title Terms: QUERY; OPTIMUM; METHOD; RELATED; DATABASE ; FINDER; JOIN; CONDITION; FORM; CHAIN; FORM; TABLE ; ACCORD; GRAPH; JOIN; THEORY

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/21 (Item 18 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

011645714 \*\*Image available\*\*  
WPI Acc No: 1998-062622/199806  
XRPX Acc No: N98-049324

GUI data access method for relational database - involves joining tables on common field using operator and defining logical schema defining table hierarchy

Patent Assignee: INFORMIX SOFTWARE INC (INFO-N)

Inventor: JACKSON B D; MALONEY C W; MAYFIELD K B; MILLS M A; TRACY K A

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 5701453	A	19971223	US 9386001	A	19930701	199806 B

Priority Applications (No Type Date): US 9386001 A 19930701

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes  
US 5701453 A 23 G06F-017/30

Abstract (Basic): US 5701453 A

The data access method involves specifying at least one logical relationship by selecting a pair of **tables** having a common field from among the relational **database tables**. The common field of the pair of **tables** is selected as a **join field**. A **join operator** is selected for the **join field** to define a **join** between the pair of **tables**, the logical relationship being one of One-to-One, One-to-Zero-or-One, One-to-Zero-or-More, and One-to-One-or-More.

A logical schema is created defining a multilevel hierarchy of the **tables** in the specified logical relationships. A first **table** and a second **table** of a pair of **tables** in a logical relationship are at an equal hierarchy level if the logical relationship is one of One-to-One and One-to-Zero-or-One. The first **table** is at a higher hierarchy level than the second **table** if the logical relationship is one of One-to-Zero-or-More One-to-One-or-More.

ADVANTAGE - Avoids user having to understand physical **database** schema to access required data.

Dwg.3/20

Title Terms: DATA; ACCESS; METHOD; RELATED; **DATABASE** ; JOIN; **TABLE** ;  
COMMON; FIELD; OPERATE; DEFINE; LOGIC; DEFINE; **TABLE** ; HIERARCHY

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

5/5/22 (Item 19 from file: 350)

DIALOG(R)File 350:Derwent WPIX

2004 Thomson Derwent. All rts. reserv.

00971226 \*\*Image available\*\*

WPI Acc No: 1992-098495/199213

XRPX Acc No: N92-073735

Computer data base and retrieval method - analysing join statements using graphical technique to determine groups of tables represented as nodes

Patent Assignee: INT BUSINESS MACHINES CORP (IBMC ); IBM CORP (IBMC )

Inventor: JACOPI T W

Number of Countries: 004 Number of Patents: 005

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
EP 476810	A	19920325	EP 91306546	A	19910718	199213 B
US 5237493	A	19940215	US 90576022	A	19900831	199407
EP 476810	A3	19931020	EP 91306546	A	19910718	199510
EP 476810	B1	19981014	EP 91306546	A	19910718	199845
DE 69130350	E	19981119	DE 630350	A	19910718	199901
			EP 91306546	A	19910718	

Priority Applications (No Type Date): US 90576022 A 19900831

Cited Patents: No-SR.Pub; 1.Jnl.Ref

Patent Details:

-----Patent No--Kind--Lan--Pg--Main--IPC-----Filing Notes-----

EP 476810 A 11

Designated States (Regional): DE FR GB

US 5237493 A 10 G06F-015/40

EP 476810 B1 E G06F-017/30

Designated States (Regional): DE FR GB

DE 69130350 E G06F-017/30 Based on patent EP 476810

Abstract (Basic): EP 476810 A

Data is retrieved using a request including a set of joint statements, each linking the name of two **tables**. The data processor includes an element to assign priority to one **table** name in each joint statement. An array of graph identifiers corresponding to the **table**

names is formed. Each group identifier is initialised to have a value representing the corresponding table name. Each joint statement is processed in succession.

The value of the priority name is substituted in the array of graph identifier values in place of each value representing the other name to derive an indication of the number of groups of linked names in the set of entered join statements.

USE/ADVANTAGE - Assures coherency of join list without need for re-entering all of elements of join list

Title Terms: COMPUTER; DATA; BASE; RETRIEVAL; METHOD; ANALYSE; JOIN;  
STATEMENT; GRAPHICAL; TECHNIQUE; DETERMINE; GROUP; TABLE ; REPRESENT;  
NODE

Derwent Class: R27; T01

International Patent Class (Main): G06F-015/40 ; G06F-017/30

International Patent Class (Additional): G06F-015/419

IPC Class: EPI

5/5/23 (Item 20 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

008743329 \*\*Image available\*\*

WPI Acc No: 1991-247345/199134

XRPX Acc No: N91-188595

Joining selected data in tables of relational data base system - by  
defining first parameter, selecting data in first table , placing in  
sub- table , defining second parameter and selecting second data

Patent Assignee: IBM CORP (IBM ); INT BUSINESS MACHINES CORP (IBM )

Inventor: CHENG J M K; HARDERLE D J; HEDGES R W; IYER B R; MOHAN C; WANG Y;  
CHENG J M; HADERLE D J

Number of Countries: 005 Number of Patents: 003

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
EP 442684	A	19910821	EP 91301085	A	19910211	199134 B
US 5241648	A	19930831	US 90479523	A	19900213	199336
JP 3104708	B2	20001030	JP 9114755	A	19910114	200057

Priority Applications (No Type Date): US 90479523 A 19900213

Cited Patents: NoSR.Pub

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes

EP 442684 A

Designated States (Regional): DE FR GB

US 5241648 A 13 G06F-015/40

JP 3104708 B2 14 G06F-017/30 Previous Publ. patent JP 4213765

Abstract (Basic): EP 442684 A

The method for joining selected data in two tables (12,10) in a relational data base system which involves defining a first parameter (35), selecting data in the first table which satisfies the first parameter and placing such selected data in order in a first sub-table (38). A second parameter is defined and data is selected in the second table which satisfies the second parameter.

Selecting data which satisfies the second parameter involves defining a preliminary parameter-based on the data in the first sub-table . Data in the second table is selected which satisfies the preliminary parameter and placed in order in a second sub table (45). Data in the second sub table which satisfies the second parameter is selected and combined with the data in the first sub table so as to join the selected data in the first and second tables (50).

ADVANTAGE - Highly efficient I/O operations. (14pp Dwg.No.5/5

Title Terms: JOIN; SELECT; DATA; TABLE ; RELATED; DATA; BASE; SYSTEM;  
DEFINE; FIRST; PARAMETER; SELECT; DATA; FIRST; TABLE ; PLACE; SUB;  
TABLE ; DEFINE; SECOND; PARAMETER; SELECT; SECOND; DATA

Derwent Class: T01

International Patent Class (Main): G06F-015/40 ; G06F-017/30

International Patent Class (Additional): G06F-012/00

File Segment: EPI

5/5/24 (Item 21 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

18:19:57 \*\*Image available\*\*  
Pub No: 1991-103669/199115  
Pub No: N91-080158

Relational data base management system for multiprocessor  
environments - allows two data base relations to be joined on  
common field in parallel relational data base field  
Patent Assignee: INT BUSINESS MACHINES CORP (IBM ); IBM CORP (IBM )  
Inventor: DIAS D M; SHI-LUNG Y P; WOLF J L; SHILUNG Y P; SHI-LUNG YU P; YU  
P S

Number of Countries: 004 Number of Patents: 005

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
EP 421408	A	19910410	EP 90119011	A	19901004	199115 B
US 5121494	A	19920609	US 89417366	A	19891005	199226
EP 421408	A3	19930421	EP 90119011	A	19901004	199401
EP 421408	B1	19970319	EP 90119011	A	19901004	199716
DE 69030228	E	19970424	DE 630228	A	19901004	199722
			EP 90119011	A	19901004	

Priority Applications (No Type Date): US 89417366 A 19891005

Cited Patents: NoSR.Pub; 3.Jnl.Ref

Patent Details:

Patent No	Kind	Lan	Pg	Main IPC	Filing Notes
-----------	------	-----	----	----------	--------------

EP 421408	A				
-----------	---	--	--	--	--

Designated States (Regional): DE FR GB

US 5121494	A		17	G06F-015/40	
------------	---	--	----	-------------	--

EP 421408	B1 E	23	G06F-017/40		
-----------	------	----	-------------	--	--

Designated States (Regional): DE FR GB

DE 69030228	E		G06F-017/40	Based on patent EP 421408	
-------------	---	--	-------------	---------------------------	--

Abstract (Basic): EP 421408 A

The join operation is performed in three stages with an optional fourth stage, the details of which depend on the underlying join algorithm used. The first stage is a preparatory stage that provides data for the second stage to use as a basis for defining the subtasks for the final join operation and to allocate subtasks evenly to different processors, even in the presence of data skew.

Once the second stage has completed its processing, the subtasks are shipped to their assigned processors for processing and the final join of the two relations in the third stage. Optionally, during the join in the third stage, there could be a dynamic reassignment of the subtasks should the join operation become unbalanced.

---

11/5/4 (Item 4 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

008461136 \*\*Image available\*\*  
WPI Acc No: 1990-348136/199046  
XRPX Acc No: N90-265999

Processing data base of system having memory unit - storing tables  
with columns and rows of data

Patent Assignee: HITACHI LTD (HITA )

Inventor: FUKUSHIMA S; KITAJIMA H; OHMSCHI K; OHSONE T; SATOH K; TANIGUCHI  
H; TSUCHIDA M; YAMAMOTO A; YAMASHITA Y

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 4967341	A	19901030	US 89418343	A	19891006	199046 B

Priority Applications (No Type Date): JP 8652438 A 19860312; JP 8628807 A  
19860214

Abstract (Basic): US 4967341 A

A set of sorted columns of one table stored in a DB machine  
is binary-searched by a hardware in synchronism with data transfer  
when another table is read from a disk unit to select a row to be  
joined.

Where a key word sequence to be searched is stored in ascending or  
descending order, determination of a max. address in an area in which  
key words smaller than a search key are stored and determination of a  
max. address in an area in which the key words smaller than or equal to  
the search key are stored are in parallel executed to determine the  
storage range in one search.

ADVANTAGE - Increases speed of join operation in relational data  
base machine.

Dwg.1/23

Title Terms: PROCESS; DATA; BASE; SYSTEM; MEMORY; UNIT; STORAGE; TABLE ;  
COLUMN; ROW; DATA

Derwent Class: T01

International Patent Class (Additional): G06F-015/20

Classification: EPI

---

14/5/4 (Item 4 from file: 347)  
DIALOG(R)File 347:JAPIO  
(c) 2004 JPO & JAPIO. All rts. reserv.

02319150 \*\*Image available\*\*  
JOIN PROCESSING SYSTEM

PUB. NO.: 62-236050 [JP 62236050 A]  
PUBLISHED: October 16, 1987 (19871016)  
INVENTOR(s): NAKAMURA SHUNICHIRO  
APPLICANT(s): MITSUBISHI ELECTRIC CORP [000601] (A Japanese Company or Corporation), JP (Japan)  
APPL. NO.: 61-079673 [JP 8679673]  
FILED: April 07, 1986 (19860407)  
INTL CLASS: [4] G06F-012/00 ; G06F-015/16  
JAPIO CLASS: 45.2 (INFORMATION PROCESSING -- Memory Units); 45.4 (INFORMATION PROCESSING -- Computer Applications)  
JOURNAL: Section: P, Section No. 685, Vol. 12, No. 104, Pg. 20, April 06, 1988 (19880406)

#### ABSTRACT

PURPOSE: To greatly shorten **join** processing time by dividing a **join** processing into plural CPU for **parallel** execution.

CONSTITUTION: The 1st **table** for execution of joins is decentralized to plural memories 7-10 and the 2nd **table** for joins is superposed on those memories 7-10 respectively. Here processors 3-6 perform the 1:1 comparison between the partial records of the 1st and 2nd **tables** and store successively the pairs of records satisfying the connecting conditions into the prescribed areas of a common memory device 1. Thus a low **table** undergone the join processing is finally produced on the device 1. As a result, the joint processing is possible in a short period without carrying out any sorting process in case one of both **tables** is especially large or small.

14/5/5 (Item 5 from file: 347)  
DIALOG(R)File 347:JAPIO  
(c) 2004 JPO & JAPIO. All rts. reserv.

01229045 \*\*Image available\*\*  
SIMULTANEOUS JOINING PROCESSING SYSTEM FOR PLURAL ITEMS

PUB. NO.: 58-166445 [JP 58166445 A]  
PUBLISHED: October 01, 1983 (19831001)  
INVENTOR(s): OHARA TOSHISAKU  
APPLICANT(s): FUJITSU LTD [000522] (A Japanese Company or Corporation), JP (Japan)  
APPL. NO.: 57-049456 [JP 8249456]  
FILED: March 27, 1982 (19820327)  
INTL CLASS: [3] G06F-007/28  
JAPIO CLASS: 45.1 (INFORMATION PROCESSING -- Arithmetic Sequence Units); 45.2 (INFORMATION PROCESSING -- Memory Units)  
JOURNAL: Section: P, Section No. 246, Vol. 08, No. 1, Pg. 97, January 06, 1984 (19840106)

#### ABSTRACT

PURPOSE: To make a line-by-line decision between acceptance and rejection through  $\leq 2$ -times matching, by utilizing the previous sequencing of data in a **table** and to exclude lines after the decision from the matching and improve processing efficiency.

CONSTITUTION: When an inquiry input for joint processing is received from a display terminal 5, relative **tables** are loaded from a data file 3 to a memory 2. Pointers A and B are initialized to indicate the uppermost lines of the **tables**. An arithmetic processing part 10 performs the matching between the **tables** A and B on the basis of retrieval conditions of shown elements. A pointer control part 11 increases the pointer A or B by one

every time the matching is carried out once to indicate the next line. Records of a coincident line are put together in output format by an input/output processing part 12 and transferred to the terminal 5.

14/5/6 (Item 1 from file: 350)

DIALOG(R) File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

015378340 \*\*Image available\*\*

WPI Acc No: 2003-439778/200341

XPX Acc No: N03-350967

Selected table joins display method in relational database management system, involves forming line between icons representing selected tables in overview diagram upon accepting selected join grid row

Patent Assignee: INT BUSINESS MACHINES CORP (IBM )

Inventor: GUTIERREZ-RIVAS H; ISMERIO F C; PAYTON B G

Number of Countries: 001 Number of Patents: 002

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 20030055830	A1	20030320	US 2001960574	A	20010920	200341 B
US 6553371	B2	20030422	US 2001960574	A	20010920	200341

Priority Applications (No Type Date): US 2001960574 A 20010920

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
US 20030055830	A1	11	G06F-007/00	
US 6553371	B2		G06F-017/30	

US 20030055830 A1

US 6553371 B2

Abstract (Basic): US 20030055830 A1

NOVELTY - A table join grid row having potential valid table join for two database tables represented in a join overview diagram, is selected from a grid. A line is formed between the icons representing the selected tables in the diagram, upon accepting the selected join grid row.

DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following:

(1) program storage device storing selected valid table joins display program; and

(2) selected valid table joins display system.

USE - For displaying selected table joins in relational database management system (RDBMS) of computer.

ADVANTAGE - The join grid and the join overview diagram are always displayed together and are synchronized. Hence, if a user chooses a join in the join grid, the join overview diagram automatically reflects the selection by drawing the line, which highlights the associated row in the join grid. Each row in the join grid automatically provides a summary of all valid two-column combinations of the selected two tables. Thus, the user need not select each of the two columns in the join grid separately. Provides efficient table join process even for a large number of tables and columns.

DESCRIPTION OF DRAWING(S) - The figure shows a flowchart explaining table joins display and selection process.

pp; 11 DwgNo 5/5

Title Terms: SELECT; TABLE; JOIN; DISPLAY; METHOD; RELATED; DATABASE; MANAGEMENT; SYSTEM; FORMING; LINE; REPRESENT; SELECT; TABLE; DIAGRAM; ACCEPT; SELECT; JOIN; GRID; ROW

Derwent Class: T01

International Patent Class (Main): G06F-007/00 ; G06F-017/30

File Segment: EPI

14/5/7 (Item 2 from file: 350)

DIALOG(R) File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

014516549 \*\*Image available\*\*

WPI Acc No: 2002-337252/200237

**Method for managing advertisement information**

Patent Assignee: VENICE SYSTEM CO LTD (VENI-N)

Inventor: YANG J H

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
KR 2001107227	A	20011207	KR 200028555	A	20000526	200237 B

Priority Applications (No Type Date): KR 200028555 A 20000526

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
KR 2001107227	A		1 G06F-017/60	

Abstract (Basic): KR 2001107227 A

NOVELTY - A method for managing advertisement information is provided to enable a client to get profit due to advertisement information, and to enable an advertiser to get more efficient effect of advertisement with less cost by using quiz game.

DETAILED DESCRIPTION - (A service operator interface part(301) interfaces communication with a computer of a service operator(100). An information classifying part(302) classifies advertisement information provided from an advertiser, quiz information provided from the service operator(100), member information inputted from a client, an input information inputted correspondingly to the quiz by the client and information about right answerer respectively. Then, the information classifying part(302) stores the classified information in corresponding database (303-307) respectively. If a client joins as member, a part of requesting input of member information(308) requests the client to input various personal information including ID and password. At the same time, the part of requesting input of member information(308) requests the client who joined as member to input the ID and the password. An external interface part(309) interfaces an operating server(300) with communication network(500). When a client registered as member connects to the operator server(300) and joins in quiz, a quiz information displaying part(310) extracts quiz information stored in a quiz information database (303), and provides the extracted quiz information to the client.

pp; 1 DwgNo 1/10

Title Terms: METHOD; MANAGE; ADVERTISE; INFORMATION

Derwent Class: T01

International Patent Class (Main): G06F-017/60

File Segment: EPI

14/5/9 (Item 4 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

014263786 \*\*Image available\*\*

WPI Acc No: 2002-084484/200212

XRPX Acc No: N02-062799

**Data retrieval optimization method for relational database management system, involves creating aggregate join index by combining commonly used columns of tables with results from computed aggregation expressions**

Patent Assignee: NCR-INT-INC (NATC-);-NCR CORP (NATC-)

Inventor: AU G K; HOANG C K; ON AU G K

Number of Countries: 027 Number of Patents: 002

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
EP 1164509	A2	20011219	EP 2001304927	A	20010606	200212 B
US 6505189	B1	20030107	US 2000594964	A	20000615	200306

Priority Applications (No Type Date): US 2000594964 A 20000615

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
EP 1164509	A2	E	12 G06F-017/30	

Designated States (Regional): AL AT BE CH CY DE DK ES FI FR GB GR IE IT

IT LU LV MC MK NL PT RO SE SI TR  
US 5505189 B1 G06F-017/30

Abstract (Basic): EP 1164509 A2

NOVELTY - The aggregation expressions from the columns of **tables** are computed. An aggregate join index is created for the **tables** by combining commonly used columns of the **tables** with results from the expressions.

DETAILED DESCRIPTION - INDEPENDENT CLAIMS are also included for the following

- (a) Computerized **database** management system;
- (b) Data structure;
- (c) Data retrieval optimization program;
- (d) Computer readable storage medium storing data retrieval optimization program

USE - For mainframe, mini computer, personal computer for use in relational **database** management system (claimed).

ADVANTAGE - Allows users to perform aggregation operations reliably using the aggregate join index. Enhances **parallel** processing across multiple access module processors.

DESCRIPTION OF DRAWING(S) - The figure shows the hardware and software environment.

pp; 12 DwgNo 1/4

Title Terms: DATA; RETRIEVAL; OPTIMUM; METHOD; RELATED; **DATABASE** ;  
MANAGEMENT; SYSTEM; AGGREGATE; JOIN; INDEX; COMBINATION; COMMON; COLUMN;  
**TABLE** ; RESULT; COMPUTATION; AGGREGATE; EXPRESS

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

14/5/10 (Item 5 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

014016959

WPI Acc No: 2001-501173/200155

XRPX Acc No: N01-371558

Exploitation of db2 universal database design rules in graphical representations

Patent Assignee: INT BUSINESS MACHINES CORP (IBM )

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
RD 439106	A	20001110	RD 2000439106	A	20001020	200155 B

Priority Applications (No Type Date): RD 2000439106 A 20001020

Patent Details:

Patent No	Kind	Lan	Pg	Main IPC	Filing Notes
RD 439106	A		2	G06F-000/00	

Abstract (Basic): RD 439106 A

NOVELTY - Relational and physical **database** design is driven by a **database** designer's heuristic knowledge of **database** design. The present invention utilizes a set of DB2 designer rules derived for DB2 Universal Database (UDB) ~~Parallel-Edition-and~~ DB2 UDB physical **database** objects.

DETAILED DESCRIPTION - These rules are useful to designers in the performance of various tasks in physical design, such as the creation of indexes to optimize the access path for minimal costs; the assignment of **tables** to appropriate tablespaces; or in the **parallel database** environment, **table** collocation, which will occur if two large sized **tables** are involved in frequent joins. Each design advice action is connected to a set of heuristic rules. Through the graphical representation of DB2 UDB design rules users can access the set of rules and tailor the set or individual rules to meet the need of their **database** applications. Examples of rule sets include activate rules, deactivate rules and modify rules. A set of heuristic rules

takes the provided data load and workload values when it generates a design proposal. If no information is available, the heuristic rule set will use predefined default values with its rule. The following figure shows Creation of Index Rules using graphical representations. The advice generated using the rule set is given in the form of a report that lists the design steps to consider. When design proposals are requested, one can either accept or reject a design proposal. The following two figures show proposed actions and a proposed report for a table object using graphical representations.

USE - None given.

pp; 2 DwgNo 0/0

Title Terms: EXPLOIT; UNIVERSAL; DATABASE ; DESIGN; RULE; GRAPHICAL;  
REPRESENT

Derwent Class: T01

International Patent Class (Main): G06F-000/00

File Segment: EPI

14/5/11 (Item 6 from file: 350)

DIALOG(R) File 350:Derwent WPIX

© 2004 Thomson Derwent. All rts. reserv.

013544726 \*\*Image available\*\*

WPI Acc No: 2001-028932/200104

XRPX Acc No: N01-022943

Edit log management system of engineering drawing, stores drawing joining management and authority registration tables based on whose contents joining of drawing data and editing authority are managed respectively  
Patent Assignee: KENSETSUISHO KENCHIKU KENYUSHOCHO (KENS-N); SHIMIZU CONSTR CO LTD (SHMC )

Number of Countries: 001 Number of Patents: 001

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
JP 2000298678	A	20001024	JP 99107319	A	19990415	200104 B

Priority Applications (No Type Date): JP 99107319 A 19990415

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
JP 2000298678	A		6 G06F-017/30	

Abstract (Basic): JP 2000298678 A

NOVELTY - The data processor (1) produces 2D engineering drawing which is split in tree form and authority of editing are managed based on contents registered in drawing joining management table (28) and authority registration table (25). The editing of drawing is performed based on parent drawing attribute data (22) and contents of updating registration table (24).

DETAILED DESCRIPTION - An INDEPENDENT CLAIM is also included for  
USE - For parallel editing and log management of engineering drawings produced during large scale construction project.

ADVANTAGE - Since editing of drawing data corresponding to parent attribute data and updated data is performed and managed based on respective tables, editing is performed simultaneously and parallelly by several operators, enabling drawing joining centralization unit to be set to the desire of user and preventing mistakes of editing same drawing by several operators.

DESCRIPTION OF DRAWING(S) - The figure shows the block diagram of edit log management system.

Data processor (1)

Parent drawing attribute data (22)

Registration tables (24,25)

Drawing joining management table (28)

pp; 6 DwgNo 1/6

Terms: EDIT; LOG; MANAGEMENT; SYSTEM; ENGINEERING; DRAW; STORAGE;  
DRAW; JOIN; MANAGEMENT; AUTHORISE; REGISTER; TABLE ; BASED; CONTENT;  
JOIN; DRAW; DATA; EDIT; AUTHORISE; RESPECTIVE

Derwent Class: T01

International Patent Class (Main): G06F-017/30

International Patent Class (Additional): G06F-017/50  
File Segment: EPI

14/5/12 (Item 7 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

013491833 \*\*Image available\*\*  
WPI Acc No: 2000-663776/200064  
Related WPI Acc No: 2003-776318  
XRPX Acc No: N00-491802

Query processing method for relational database management system,  
involves performing join operation of data streams locally, after  
detecting appropriate conjunct predicates

Patent Assignee: INT BUSINESS MACHINES CORP (IBMC )  
Inventor: LOHMAN G M; PIRAHESH M H; SHEKITA E J; SIMMEN D E; URATA M S  
Number of Countries: 001 Number of Patents: 001  
Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 6112198	A	20000829	US 9751259	P	19970630	200064 B
			US 98106473	A	19980629	

Priority Applications (No Type Date): US 9751259 P 19970630; US 98106473 A  
19980629

Patent Details:

Patent No	Kind	Lan Pg	Main IPC	Filing Notes
6112198	A	25	G06F-017/30	Provisional application US 9751259

Abstract (Basic): US 6112198 A

NOVELTY - Two partitioned data streams that relate to join  
operation, are received. Usage possibility of conjunct predicates to  
locally perform a parallel inner join or outer join operation, is  
determined. The join operation is performed locally, if the detection  
result approves the usage of appropriate conjunct predicates.

USE - For parallel query processing in SQL processing relational  
database managements system (RDBMS) in computer network.

ADVANTAGE - Enables to optimize or avoid data repartitioning by  
recognizing the possible partitioning requirements for achieving  
parallelism for a query operation, and when the partitioning property  
of data satisfies the partitioning requirements of query operation.

DESCRIPTION OF DRAWING(S) - The figure represents the query  
execution plan for directed join operation in DBMS.

pp: 25 DwgNo 4/12

Title Terms: QUERY; PROCESS; METHOD; RELATED; DATABASE ; MANAGEMENT;  
SYSTEM; PERFORMANCE; JOIN; OPERATE; DATA; STREAM; LOCAL; AFTER; DETECT;  
APPROPRIATE

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

14/5/13 (Item 8 from file: 350)  
DIALOG(R)File 350:Derwent WPIX  
(c) 2004 Thomson Derwent. All rts. reserv.

013491833 \*\*Image available\*\*  
WPI Acc No: 1999-214411/199918  
XPIX Acc No: N99-157814

Proximity join operations performing method on high dimensional delta  
points in microprocessor system

Patent Assignee: INT BUSINESS MACHINES CORP (IBMC )  
Inventor: AGRAWAL R; SHAFER J C  
Number of Countries: 001 Number of Patents: 001  
Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
US 5884320	A	19990316	US 97920331	A	19970820	199918 B

Priority Applications (No Type Date): US 97920331 A 19970820

Patent Details:

Patent No Kind Lan Pg Main IPC Filing Notes  
US 5884320 A 17 G06F-017/30

Abstract (Basic): US 5884320 A

NOVELTY - The join operations are assigned to the processors using index structures. The data points are simultaneously redistributed and joined in the processors in parallel based on predetermined joining conditions.

DETAILED DESCRIPTION - The data points among the processors are partitioned and an index structure for these data points is created. The index structure has several leaf nodes corresponding to a subset of data points. INDEPENDENT CLAIMS are also included for the following:

- (a) computer product;
- (b) database system.

USE - In microprocessor system.

ADVANTAGE - Builds index structure and performs joins efficiently using minimum amount of storage space due to system's parallelism.

DESCRIPTION OF DRAWING(S) - The figure shows flowchart of overall operational sequence of performing join operations on data points of database in parallel.

pp: 17 DwgNo 3/10

Title Terms: PROXIMITY; JOIN; OPERATE; PERFORMANCE; METHOD; HIGH; DIMENSION  
; DELTA; POINT; MICROPROCESSOR; SYSTEM

Derwent Class: T01

International Patent Class (Main): G06F-017/30

File Segment: EPI

14/5/14 (Item 9 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

010421383 \*\*Image available\*\*

WPI Acc No: 1995-322699/199542

XRPX Acc No: N95-242922

Database join processing system - joins relations based on join fields in relational database

Assignee: MITSUBISHI DENKI KK (MITQ ); MITSUBISHI ELECTRIC CORP

Inventor: MATSUMOTO T

Number of Countries: 003 Number of Patents: 004

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
GB 2287807	A	19950927	GB 952768	A	19950213	199542 B
JP 7253991	A	19951003	JP 9445620	A	19940316	199548
US 5613142	A	19970318	US 95388616	A	19950214	199717
GB 2287807	B	19980506	GB 952768	A	19950213	199820

Priority Applications (No Type Date): JP 9445620 A 19940316

Patent Details:

Patent No	Kind	Lan	Pg	Main IPC	Filing Notes
GB 2287807	A	124	G06F-007/32		
JP 7253991	A	24	G06F-017/30		
US 5613142	A	48	G06F-007/08		
GB 2287807	B		G06F-007/32		

Abstract (Basic): GB 2287807 A

The system joins distributed data with a join key and produces a joined table (100,200). Recording devices, disk drives (4a-4d) store the distributed data e.g. employee data (5a-5d) and sales data (6a-6d). Slave-processors (3a-3d) are coupled to the recording devices to retrieve the data and output the data. The main processor (1) receives the data from the slave-processors and produces the joined table.

Each slave-processor checks a join key of the second data, sales data, with a join key of the first data. Based on the checking result it selects the second data (400a-400d) and outputs the data to the main processor.

USE/ADVANTAGE - Provides high speed joining , eliminates burden on master processor as data distributed and stored in slave processors, part of join processing done in parallel .

Dwg.3/23

CLASS: DATABASE ; JOIN; PROCESS; SYSTEM; JOIN; RELATED; BASED; JOIN; RELATED; DATABASE

CLASS: T01

International Patent Class (Main): G06F-007/08 ; G06F-007/32 ; G06F-017/30

International Patent Class (Additional): G06F-007/14 ; G06F-007/36 ; G06F-012/00 ; G06F-012/04

File Segment: EPI

14/5/16 (Item 11 from file: 350)

DIALOG(R)File 350:Derwent WPIX

(c) 2004 Thomson Derwent. All rts. reserv.

000571767 \*\*Image available\*\*

WPI Acc No: 1991-075800/199111

XRFX Acc No: N91-058589

Join processor for relational database - extracts join fields of relations from aux. memory and sorts with aux. processors, dividing and storing in parallel

Patent Assignee: MITSUBISHI DENKI KK (MITQ )

Inventor: MINEMURA H; NAKAMURA S

Number of Countries: 003 Number of Patents: 004

Patent Family:

Patent No	Kind	Date	Applicat No	Kind	Date	Week
GB 2235798	A	19910313	GB 9018947	A	19900830	199111 B
GB 2235798	A	19910301				199119
GB 2235798	B	19930901	GB 9018947	A	19900830	199335
US 5247662	A	19930921	US 90576202	A	19900829	199339

Priority Applications (No Type Date): JP 90151114 A 19900608; JP 89225815 A 19890831

Patent Details:

Patent No	Kind	Lan	Pg	Main IPC	Filing Notes
GB 2235798	B		20	G06F-015/40	
US 5247662	A		13	G06F-007/24	

Abstract (Basic): GB 2235798 A

A join processor for relational **databases** includes a main processor (1) for processing information. A main memory (2) is connected to the main processor to store information. A number of auxiliary processors (3) are controlled by the main processor.

A number of auxiliary memories (4) are each connected to the corresp. auxiliary processors. A device enables the auxiliary processors to perform in parallel dividing and storing relations of relational **databases** in the auxiliary memories in units of record and extracting from the auxiliary memories and sorting join fields of the relations to be joined.

USE - Fast join processor for a large size of relations to be joined, e.g., employee records. (21pp Dwg.No.1/6

---

File 348:EUROPEAN PATENTS 1978-2004/Apr W02

(c) 2004 European Patent Office

File 349:PCT FULLTEXT 1979-2002/UB=20040415,UT=20040408

(c) 2004 WIPO/Univentio

Set	Items	Description
S1	7315	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (7N) (ROW? ? OR TUPLE? ? OR COLUMN? ? OR FIELD? ? OR CELL? ? OR TABLE? ?)
S2	1422	S1(7N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIME OR AS OR WHILE)
S3	549024	DATABASE? ? OR DATA()BASE? ? OR TABLE? ?
S4	13343	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (20N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIME)
S5	10265	REDISTRIBUT? OR RE()DISTRIBUT?
S6	5	S2(50N)S3(50N)S5
S7	0	S3(50N)S4(50N)S5
S8	84186	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (20N) (WHILE OR AS)
S9	6	S3(50N)S5(50N)S8
S10	317246	ROW? ? OR TUPLE? ? OR COLUMN? ?
S11	18855	S10(5N) (REDISTRIBUT? OR DISTRIBUT? OR SHUFFL? OR RESHUFFL? OR SHIFT??? OR RESHIFT??? OR MOV??? OR MIX??? OR REORDER??? OR SORT??? OR RESORT??? OR REARRANG? OR REORGANI? OR REGROUP??? OR RE() (ARRANG? OR ORGANI? OR GROUP??? OR ORDER?))
S12	3266	S1(25N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC - OR SAME()TIME OR AS OR WHILE)
S13	15	S12(50N)S11(50N)S3
S14	199	S2(50N)S3 AND IC=G06F
S15	10	S2(50N)S3(50N)S11 AND IC=G06F
S16	4	S15 NOT S13
S17	1238781	DATA OR INFORMATION OR FIELD? ? OR CELL? ?
S18	149099	S17(5N) (REDISTRIBUT? OR DISTRIBUT? OR SHUFFL? OR RESHUFFL? OR SHIFT??? OR RESHIFT??? OR MOV??? OR MIX??? OR REORDER??? OR SORT??? OR RESORT??? OR REARRANG? OR REORGANI? OR REGROUP??? OR RE() (ARRANG? OR ORGANI? OR GROUP??? OR ORDER?))
S19	33	S12(50N)S18(50N)S3 AND IC=G06F
S20	24	S2(50N)S3(50N)S18 AND IC=G06F
S21	34	S19:S20
S22	29	S21 NOT (S6 OR S9 OR S13 OR S16)

---

13/3,K/1 (Item 1 from file: 348)  
DIALOG(R)File 348:EUROPEAN PATENTS  
(c) 2004 European Patent Office. All rts. reserv.

01713701

Eliminating group-by operation in a join plan

Beseitigen von "group-by"-Befehlen aus einem Verbindungsplan

Elimination des operations d'aggregation "group-by" dans un plan de jointure

PATENT ASSIGNEE:

NCR INTERNATIONAL INC., (1449480), 1700 South Patterson Boulevard,  
Dayton, Ohio 45479, (US), (Applicant designated States: all)

INVENTOR:

Pham, Son, 18100 Herbold Street, Northridge, CA 91325, (US)

Pham, Thu K., 18100 Herbold Street, Northridge, CA 91325, (US)

LEGAL REPRESENTATIVE:

Williamson, Brian et al (84715), International IP Department, NCR  
Limited, 206 Marylebone Road, London NW1 6LY, (GB)

PATENT (CC, No, Kind, Date): EP 1403788 A2 040331 (Basic)

APPLICATION (CC, No, Date): EP 2003255226 030823;

PRIORITY (CC, No, Date): US 259070 020927

DESIGNATED STATES: AT; BE; BG; CH; CY; CZ; DE; DK; EE; ES; FI; FR; GB; GR;  
HU; IE; IT; LI; LU; MC; NL; PT; RO; SE; SI; SK; TR

EXTENDED DESIGNATED STATES: AL; LT; LV; MK

INTERNATIONAL PATENT CLASS: G06F-017/30

ABSTRACT WORD COUNT: 62

NOTE:

Figure number on first page: 1

LANGUAGE (Publication,Procedural,Application): English; English; English

FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	200414	648
SPEC A	(English)	200414	11609
Total word count - document A			12257
Total word count - document B			0
Total word count - documents A + B			12257

...SPECIFICATION that such attributes were at some point part of a candidate key (since they were originally associated with data elements having Value(underscore)1 that have subsequently been changed to Value(underscore)3 as a result of a join operation). As explained below, this change from Value(underscore)1 to Value(underscore)3 was made due to the fact that the join may cause duplicate values to occur in the result, thereby rendering the candidate key to no longer be unique. However, after the group by, duplicates are removed and the candidate key is again unique for each row.

If neither condition 512 nor 516 is satisfied, then the optimizer module 20 sets ( at 520) the data elements of all attributes (the grouping fields) to Value(underscore)2. In effect, as part of the partial group -by operation on a table T, the grouping fields are set to either Value(underscore)1 or Value(underscore)2, depending on the conditions noted above.

The optimizer module 20 also updates (at 522) the data structure 401 in response to a join of two tables (table T1 and...

13/3,K/2 (Item 2 from file: 348)  
DIALOG(R)File 348:EUROPEAN PATENTS  
(c) 2004 European Patent Office. All rts. reserv.

01566388

Partitioned database system

Partitioniertes Datenbanksystem

Systeme partitionne de base de donnees

PATENT ASSIGNEE:

NCR INTERNATIONAL INC., (1449480), 1700 South Patterson Boulevard,  
Dayton, Ohio 45479, (US), (Applicant designated States: all)

INVENTOR:

Sinclair, Paul L., 1825 John St., Manhattan Beach, CA 90266, (US)  
Cohen, Steven B., 1706 Haynes Lane, Redondo Beach, CA 90278, (US)  
Pederson, Donald R., 12410 Pathos Lane, San Diego, CA 92129, (US)

LEGAL REPRESENTATIVE:

Cleary, Fidelma et al (85871), International IP Department NCR Limited  
206 Marylebone Road, London NW1 6LY, (GB)

PATENT (CC, No, Kind, Date): EP 1302873 A2 030416 (Basic)

APPLICATION (CC, No, Date): EP 2002256960 021008;

PRIORITY (CC, No, Date): US 981613 011016

DESIGNATED STATES: AT; BE; BG; CH; CY; CZ; DE; DK; EE; ES; FI; FR; GB; GR;  
IE; IT; LI; LU; MC; NL; PT; SE; SK; TR

EXTENDED DESIGNATED STATES: AL; LT; LV; MK; RO; SI

INTERNATIONAL PATENT CLASS: G06F-017/30

ABSTRACT WORD COUNT: 144

NOTE:

Figure number on first page: 2

LANGUAGE (Publication,Procedural,Application): English; English; English

FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	200316	597
SPEC A	(English)	200316	3752
Word count - document A			4349
Word count - document B			0
Total word count - documents A + B			4349

...SPECIFICATION facilities 1201-4)) by the parsing engine 130 (not shown).  
For example, two columns 210, 220 can be designated as the primary index  
when the **table** is created. The hash function is then applied to the  
contents of columns 210, 220 for each row. The hash bucket portion of the  
resulting...

...number and the hash function is the sum of remainder when the row number  
is divided by four and the value one, the first eight **rows** will be  
**distributed** as shown in Fig. 2.

Queries involving the values of columns in the primary index can be  
efficiently executed because the processing module 110n)) having...

...from row 2 are desired, the parsing engine 130 can apply the hashing  
function to determine that only processing module 1102)) need to be used.

As another example, an equality join between two **tables** that have the  
same primary index columns is very efficient. All of the **rows** that need  
to be **joined** are found in the same data storage facility 120n)) and no  
movement of information from rows between the facilities is necessary.

While the primary index of a **table** can be chosen for equality  
joins, for example the order number **column** of an order **table**,  
additional design features can make range searches, for example a range  
of dates from the date column, more efficient. Referring to Fig. 3, a  
...tioned...

13/3,K/3 (Item 3 from file: 348)

DIALOG(R)File 348:EUROPEAN PATENTS

(c) 2004 European Patent Office. All rts. reserv.

01561046

Providing a join plan using group-by operator

Bereitstellen eines join-Planes unter Benutzung des group-by-Operators

Provision d'un plan de raccordement en utilisant de l'opérateur "group-by"

PATENT ASSIGNEE:

NCR International, Inc., (1449484), 1700 South Patterson Boulevard,  
Dayton, Ohio 45479, (US), (Applicant designated States: all)

INVENTOR:

Pham, Son, 18100 Herbold Street, Northridge, CA 91325, (US)

Pham, Thu K., 18100 Herbold Street, Northridge, CA 91325, (US)

LEGAL REPRESENTATIVE:

Cleary, Fidelma et al (85871), International IP Department NCR Limited

206 Marylebone Road, London NW1 6LY, (GB)  
PATENT (CC, No, Kind, Date): EP 1298543 A2 030402 (Basic)  
APPLICATION (CC, No, Date): EP 2002256224 020909;  
PRIORITY (CC, No, Date): US 967561 010928  
DESIGNATED STATES: AT; BE; BG; CH; CY; CZ; DE; DK; EE; ES; FI; FR; GB; GR;  
IE; IT; LI; LU; MC; NL; PT; SE; SK; TR  
EXTENDED DESIGNATED STATES: AL; LT; LV; MK; RO; SI  
INTERNATIONAL PATENT CLASS: G06F-017/30  
ABSTRACT WORD COUNT: 89  
NOTE:

Figure number on first page: 2

LANGUAGE (Publication,Procedural,Application): English; English; English  
FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	200314	1475
SPEC A	(English)	200314	10398
Total word count - document A			11873
Total word count - document B			0
Total word count - documents A + B			11873

...SPECIFICATION the further Group-by on A of the Group-by on B has an identical result of group-by A. Mathematically, if  $t$  is a **table**, then Group-by on B ( $t$ ) (contains subset) Group-by on A( $t$ ).  
Group-by on A (Group-by on B ( $t$ )) = Group-by on A( $t$ ).  
The Group-by on B groups or **sorts** all **rows** using attributes in B before an aggregation is performed. Since B is larger than A, the set of rows with a ... is the subset of rows with the same constant in A. Hence, group-by B is a finer partition of group-by A.  
Given two **tables**  $t_1$  and  $t_2$ , the following proposition (referred to as "Proposition 1") is correct:  
 $(t_1 \times t_2)' = (t_1' \times t_2')'$   
 $(t_1 \times t_2)' = (t_1' \times t_2)'$   
This provides flexibility in how the **tables** are **joined** to achieve the final **join** result. The optimizer module 20 can thus choose among plural join paths that involve Groupby **tables**.  
Given a **table**  $t$ , a Group-by on attribute  $x$  is a partition of the rows into the classes of the same value in  $x$ . If the table...

13/3,K/4 (Item 4 from file: 348)  
DIALOG(R)File 348:EUROPEAN PATENTS  
(c) 2004 European Patent Office. All rts. reserv.

01239309

METHOD FOR COMBINING TABLE DATA  
VERFAHREN ZUM KOMBINIEREN VON TABELLEN-DATEN  
PROCEDE DE COMBINAISON DE DONNEES DE TABLEAU  
PATENT ASSIGNEE:

Turbo Data Laboratory Inc., (3191480), Court house Kikuna 804, 1101-7,  
Matsumi-cho 4-chome, Kanagawa-ku, Yokomaha-shi, Kanagawa 221-0005, (JP)  
, (Applicant designated States: all)

INVENTOR:

FURUSHO, Shinji, Court house Kikuna 804, 1101-7, Matsumi-cho 4-chome,  
Kanagawa-ku, Yokohama-shi, Kanagawa 221-0005, (JP)

-----LEGAL REPRESENTATIVE:-----

Zimmermann, Gerd Heinrich (78963), Zimmermann & Partner, P.O. Box 33 09  
20, 80069 Munchen, (DE)

PATENT (CC, No, Kind, Date): EP 1191462 A1 020327 (Basic)  
WO 200073939 001207

APPLICATION (CC, No, Date): EP 2000929916 000530; WO 2000JP3465 000530

PRIORITY (CC, No, Date): JP 99151156 990531

DESIGNATED STATES: AT; BE; CH; CY; DE; DK; ES; FI; FR; GB; GR; IE; IT; LI;  
LU; MC; NL; PT; SE

EXTENDED DESIGNATED STATES: AL; LT; LV; MK; RO; SI

INTERNATIONAL PATENT CLASS: G06F-017/30; G06F-019/00

ABSTRACT WORD COUNT: 239

NOTE:

Page number on first page: 11

LANGUAGE (Publication,Procedural,Application): English; English; Japanese  
FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	200213	6945
SPEC A	(English)	200213	18326
Total word count - document A			25271
Total word count - document B			0
Total word count - documents A + B			25271

...SPECIFICATION array (see symbol 2401 of Fig. 24), the number of record numbers starting from what position that are associated with record numbers in the master table can be found from the count and start position, respectively. To wit, the aforementioned position matches the position indicated by the count and the aforementioned...

...total and start position looked up in Step 2703, and in the array of pointers to the value list within the information block including the fields to be displayed in the joined table (view), the pointer values indicating the various record numbers are fetched and rearranged in order as a new pointer array (Step 2703). Thereby, among the group of information blocks to be joined, those to be derived from the master table are complete upon their creation.

In Fig. 28, regarding record number "0," "0" is the pointer value to the corresponding position (row 1) within the...

...are fetched. From these values, one can see that the "2-0=2" record numbers "1" and "3" from the position corresponding to "0" (row 1) in the sorted set (sort array) are the record numbers of the slave table associated to record number "0" of the master table. Accordingly, on the slave table side, the pointer values of "0" and "1" are fetched sequentially from the positions indicated by the record numbers "1" and "3" (row 2 and row 4) within the information block for the "Game date" field to be displayed in the table (view), and these are placed as new pointers in the value list of the information block for "Ticket/Game date" which is an information block for the joined table (view). Note that the value list contained in this information block for "Ticket/Game date" is identical to the value list contained in the information...

13/3,K/5 (Item 5 from file: 348)  
DIALOG(R)File 348:EUROPEAN PATENTS  
(c) 2004 European Patent Office. All rts. reserv.

00480464

Joining selected data in tables of a relational data base system.

Verbinden von ausgewählten Daten in Tabellen eines relationellen Datenbanksystems.

Joindre des données sélectionnées de tableaux d'un système de base de données relationnelles.

PATENT ASSIGNEE:

International Business Machines Corporation, (200120), Old Orchard Road,  
Armonk, N.Y. 10504, (US), (applicant designated states: DE;FR;GB)

INVENTOR:

-----Cheng, Josephine Miu-Kung, 1248 Valley Quail Circle, San Jose, California  
95120, (US)

Haderle, Donald James, 812 Lilac Way, Los Gatos, California 95032, (US)

Hedges, Richard William, 98 Pegasus Trail, Scarborough, Ontario, M1G 3P1,  
(CA)

Iyer, Balakrishna Raghavendra, 39503 Ross Common, Fremont, California  
94538, (US)

Mohan, Chandrasekaran, 727 Portswood Drive, San Jose, California 95120,  
(US)

Wang, Yun, 766 Prestwick Court, Sunnyvale, California 94087, (US)

LEGAL REPRESENTATIVE:

Atchley, Martin John Waldegrave (27831), IBM United Kingdom Limited  
Intellectual Property Department Hursley Park, Winchester Hampshire

SO21 2JN, (GB)  
PATENT (CC, No, Kind, Date): EP 442684 A2 910821 (Basic)  
APPLICATION (CC, No, Date): EP 91301085 910211;  
PRIORITY (CC, No, Date): US 479523 900213  
DESIGNATED STATES: DE; FR; GB  
INTERNATIONAL PATENT CLASS: G06F-015/40;  
ABSTRACT WORD COUNT: 181

LANGUAGE (Publication,Procedural,Application): English; English; English  
FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	EPABF1	355
SPEC A	(English)	EPABF1	6342
Total word count - document A			6697
Total word count - document B			0
Total word count - documents A + B			6697

...SPECIFICATION of the outer table for rows which satisfy the local condition; then, for each of these rows, a complete join column scan of the inner Table 10 or a probe via an index on the join column is made to find matching records.

The sort/merge join of the prior art is illustrated in Figure 3 of the accompanying drawings. A pre-condition of this technique is that the outer and inner tables be ordered based upon their join columns, which enables the join to be done similarly to the MERGE phase of a SORT routine. Relatedly, a table can be ordered by sorting, or reached through an index to provide access in join column sequence. Local predicates on each table are applied before a sort, as illustrated in Figure 3. Thus, the Table 20 is built from the Table 12 by sorting the manager = Davis rows in order by employee number magnitude. Similarly, a sorted Table 22 is derived from the skill Table 10 by selecting skill table records having "test" in their skill fields and ordering those records by employee number. Once join column order is imposed on the tables, the join is done very efficiently using the join column sorted structures. Thus, the join begins by scanning the join column of the Table 22 with the employee number 53, corresponding to the first entry in...

...and the scan of the join column of the Table 22 is picked up from where the last scan stopped. In this manner, the join columns of the sorted Tables 20 and 22 are each scanned only once in effecting the join of the tables. In comparison, the nested loop join procedure requires that...

...table row is found satisfying the outer table local predicate or predicates.

The nested loop join technique makes efficient use of an index on the join column of the inner table. The nested loop join technique is used when the join column values passed to the inner table are in sequence and the join column index of the inner table is clustered, and when the number of rows retrieved in the inner table by finding matching values in the index is small. As is known, an index is clustered when the rows of the table are, for the most part, stored in the same physical sequence as the sequence of key values.

The drawbacks of the nested loop join are...scan on T2 is open, Then Close it.

If scan on intermediate result table is open, then

~~Close it.~~

The scan on the index of Table 10 is discussed in greater detail below. This scan progresses sequentially through the index on the join column of the Table 10 to retrieve the...

...RID's of the Table 10 rows in an index format having the form (key value:RID list). Since the scan is keying on the join column of Table 10, the key value is the value of the join column of an outer table row, while the parameter "RID list" is a list which identifies table rows which have that key value.

The procedure 35 in Figure 4 sorts the rows of Table 12 by the local predicate or parameter (manager = "Davis") and orders the corresponding records by their join column values. The outcome of the

process is represented by the Table 38.

Returning to the open cursor process of Flow Chart I, a DO loop is opened in step 102 which is continued until the end of file (EOF) indicators are encountered in Table 12. In the procedure, for each of the rows in the outer table, Table 12, which satisfies the predicates or parameters which are local to the outer table, the RID list is obtained from the index on the Table...illustrates the embodiment of the invention described above and embodied in Flow Charts I-III. One principal precondition of the embodiment is that the outer table be ordered by sorting or indexing on the join predicates or table space scan with a well-clustered index on the join column. In Figure...

...is not limited to ordering the outer table by the procedure 35. As is known, if an index exists on one or more of the join columns of the outer table, ordering is implicit in the index.

The technique described also presupposes ordering of the inner table by indexing on its join columns. Again, indexing implies ordering and supports sequential access to the rows of the table if the RID's are ordered.

Thus, using the rows of the outer Table 12, ordered and sorted by the join column, the technique described takes, as an index key value, the employee number of the first record, which is 53. This key value is used to scan the index of Table 10. In this regard, the scan on the index of Table 10 starts at the set (1:9) and scans from there down to the set (53:1,11). The list (1,11) is fetched from the set and two composite rows are built and placed in the Table 45. When placed in the table, the rows are ordered, so that the first row is that denoted by the RID value of 1 and the second by the RID value...

13/3,K/6 (Item 6 from file: 348)  
DIALOG(R)File 348:EUROPEAN PATENTS  
(c) 2004 European Patent Office. All rts. reserv.

00421295

Joining two database relations on a common field in a parallel relational database field

Verfahren zur Verbindung zweier Relationen einer Datenbank auf einem gemeinsamen Feld in einem parallelen Datenbanksystem

Methode pour joindre deux relations d'une base de donnees sur un champ commun dans un systeme de base de donnees parallele

PATENT ASSIGNEE:

International Business Machines Corporation, (200120), Old Orchard Road,  
Armonk, N.Y. 10504, (US), (applicant designated states: DE;FR;GB)

INVENTOR:

Dias, Daniel Manuel, 16 Pike Place, Mahopac, New York 10541, (US)  
Wolf, Joel Leonard, Cherokee Court, Goldens Bridge, New York 10526, (US)  
Shi-Lung Yu, Philip, 18 Stornoway, Chappaqua, New York 10514, (US)

LEGAL REPRESENTATIVE:

Schafer, Wolfgang, Dipl.-Ing. et al (62021), IBM Deutschland  
Informationssysteme GmbH Patentwesen und Urheberrecht, 70548 Stuttgart,  
(DE)

PATENT (CC, No, Kind, Date): EP 421408 A2 910410 (Basic)  
EP 421408 A3 930421  
EP 421408 B1 970319

EP 90119011 901004;

PATENT (CC, No, Date): US 417366 891005

DESIGNATED STATES: DE; FR; GB

INTERNATIONAL PATENT CLASS: G06F-017/40;

ABSTRACT WORD COUNT: 220

LANGUAGE (Publication,Procedural,Application): English; English; English

FULLTEXT AVAILABILITY:

Available Text	Language	Update	Word Count
CLAIMS A	(English)	EPABF1	1183
SPEC A	(English)	EPABF1	7626
Total word count - document A			8809
Total word count - document B			0

...SPECIFICATION data skew by partitioning the join operation into separate jobs and optimally scheduling the jobs among a plurality of processors.

A common operation in relational database systems is the natural join of two relations on respective columns defined over a common domain. See, for example, the description of the natural join...

...row is the concatenation of two rows, one from each of the original relations, such that both rows have the same value in their respective join columns.

One popular algorithm for computing the join of two relations is the sort-merge technique as described by M. Blasgen and K. Eswaran in "Storage and Access in Relational Databases", IBM Systems Journal, vol. 4, pp. 363 et seq. (1977). It can be summarized briefly as follows: First, each of the relations is sorted (if necessary) according to the join column. Second, the two sorted relations are scanned in the obvious interlocked sequence and merged for rows which have equal values.

When sort-merge joins are performed in parallel on a multiple processor database system, there exists a problem of data skew that might exist in the join columns of the relations. In general, the issue of skew is not addressed by the join algorithms described in the literature. An early article on parallel...

...runs from disk. The merge tree is mapped to different processors with the final merge being sequential.

In "Join and Semijoin Algorithms for a Multiprocessor Database Machine", ACM Trans. on Database Machines, vol. 9, no. 1, March 1984, pp. 133-161, P. Valduriez and G. Gardarin describe the algorithm generalized to...

13/3,K/7 (Item 1 from file: 349)

DIALOG(R)File 349:PCT FULLTEXT

(c) 2004 WIPO/Univentio. All rts. reserv.

00986561 \*\*Image available\*\*

USING ASSOCIATIVE MEMORY TO PERFORM DATABASE OPERATIONS

UTILISATION D'UNE MEMOIRE ASSOCIATIVE POUR EXECUTER DES OPERATIONS DE BASE DE DONNEES

Patent Applicant/Assignee:

ETAGON LIMITED, Euro-American Buildings, R.G. Hodge Plaza Wickhams Cay 1, P.O. Box 3161, Road Town, Tortola, VG, GB (Residence), GB (Nationality), (For all designated states except: US)

Patent Applicant/Inventor:

ZAROM Rony, 181 E 65th Street, 14th Floor, New York, NY 10021, US, US (Residence), IL (Nationality), (Designated only for: US)

Yeh Kenneth, 518 West 111th Street # 64, New York, NY 10025, US, US (Residence), US (Nationality), (Designated only for: US)

Yeh Kenneth, 4030 75th Street, Apt. 7G, Elmhurst, NY 11373, US, US (Residence), US (Nationality), (Designated only for: US)

Legal Representative:

G E EHRLICH (1995) LTD (agent), 28 Bezalel Street, 52 521 Ramat Gan, IL,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200317136-A1-20030227- (WO-0317136)

Application: WO 2002IL677 20020815 (PCT/WO IL0200677)

Priority Application: US 2001312778 20010816

Designated States: AE AG AL AM AT (utility model) AT AU AZ BA BB BG BR BY BZ CA CH CN CO CR CU CZ (utility model) CZ DE (utility model) DE DK (utility model) DK DM DZ EC EE (utility model) EE ES FI (utility model) FI GB GD GE GH GM HR HU ID IL IN IS JP KE KG KP KR KZ LC LK LR LS LT LU LV MA MD MG MK MN MW MX MZ NO NZ OM PH PL PT RO RU SD SE SG SI SK (utility model) SK SL TJ TM TN TR TT TZ UA UG US UZ VC VN YU ZA ZM ZW (EP) AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LU MC NL PT SE SK TR (OA) BF BJ CF CG CI CM GA GN GQ GW ML MR NE SN TD TG (AP) GH GM KE LS MW MZ SD SL SZ TZ UG ZM ZW (EA) AM AZ BY KG KZ MD RU TJ TM

Publication Language: English  
Filing Language: English  
Fulltext Word Count: 11058

Fulltext Availability:  
Detailed Description

#### Detailed Description

... can be a table scan, a full index scan or a partial index scan for example.

For a join statement that joins more than 2 tables, the query optimizer chooses which pair of tables is joined first and then which table is joined to the result, and so on. The query optimizer then chooses an operation to use to perform the join operation.

In a row set join, the inner row set is accessed to find all the matching rows. Therefore, in a nested loop join, the inner row set is accessed as many times as the number of rows in the outer row set.

In a sort merge join, the two row sets being joined are sorted by the join keys if they are not already in key order.

In a hash join, the inner row set is hashed into memory, and a hash table is built using the join key, which is the probe key for the join operation. Each row from the outer row set is then hashed, and the hash table is probed to join all matching rows. If the inner row set is very large, then only a portion of it is hashed into memory...the optimizer decides

that the amount of data is large enough to warrant a hash join, or it is unable to drive from the outer table to the inner table. The outer table (with preserved rows) is used to build the hash table, and the inner table is used to probe the hash table.

Sort merge joins can be used to join rows from two independent sources.

Sort merge joins are useful when the join condition between two tables is an inequality condition (but not a nonequality) like <, <=, >, or >=. In a merge join, there is no concept of a driving table. This type of join operation may optionally be performed as follows.

1. Sort join operation: Both inputs are sorted on the join key.

2. Merge join operation: The sorted lists are merged together.

If the input is already sorted by the join column, then a sort join operation is not performed for that row source.

Sort merge outerjoins are used when an outerjoin cannot drive from the outer table to the inner table.

A

13/3,K/8 (Item 2 from file:- 349)-  
DIALOG(R)File 349:PCT FULLTEXT  
(c) 2004 WIPO/Univentio. All rts. reserv.

09/06/08 \*\*Image available\*\*

SYNCHRONOUS CHANGE DATA CAPTURE IN A RELATIONAL DATABASE  
CAPTURE DE DONNEES DE CHANGEMENT SYNCHRONE DANS UNE BASE DE DONNEES  
RELATIONNELLE

Patent Applicant/Assignee:

ORACLE INTERNATIONAL CORPORATION, 500 Oracle Parkway, MS 50P7, Redwood  
Shores, CA 94065, US, US (Residence), US (Nationality), (For all  
designated states except: US)

Patent Applicant/Inventor:

DELOTT William D, 18 Powers Road, Hollis, NH 03049, US, US (Residence),  
US (Nationality), (Designated only for: US)

Legal Representative:

CARLSON Stephen C (et al) (agent), Ditthavong & Carlson, P.C., 10507  
Braddock Rd, Suite A, Fairfax, VA 22032, US,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200295632 A2-A3 20021128 (WO 0295632)

Application: WO 2002US16470 20020524 (PCT/WO US0216470)

Priority Application: US 2001863422 20010524

Designated States: AE AG AL AM AT AU AZ BA BB BG BR BY BZ CA CH CN CO CR CZ

DE DK DM DZ EC EE ES FI GB GD GE GH GM HR HU ID IL IN IS JP KE KG KR KZ

LC LK LR LS LT LU LV MA MD MG MK MN MW MX MZ NO NZ OM PH PL PT RO RU SD

SE SG SI SK SL TJ TM TR TT TZ UA UG US UZ VN YU ZA ZM ZW

(EP) AT BE CH CY DE DK ES FI FR GB GR IE IT LU MC NL PT SE TR

(OA) BF BJ CF CG CI CM GA GN GQ GW ML MR NE SN TD TG

(AP) GH GM KE LS MW MZ SD SL SZ TZ UG ZM ZW

(EA) AM AZ BY KG KZ MD RU TJ TM

Publication Language: English

Filing Language: English

Fulltext Word Count: 6995

Fulltext Availability:

Detailed Description

Detailed Description

... To obtain the contents of the change table 223 in the order in which  
the operation originally occurred, a database join operation between the  
change table 223 and the transaction 240 can be used. For example, one  
implementation of the present invention may use the following SQL join  
operation, in which change table 223 has source table columns 231  
named C1, ...that all SQL statements within a given  
transaction have the same value for a transaction identifier and that  
this value was stored in the transaction table 240 in the same row as  
the commit system change number 243. Therefore a join across the two  
tables matches up all SQL statements belonging to a particular  
transaction with their associated commit system change  
number 243. Moreover, the SQL ORDER BY clause returns the change rows  
in  
increasing sorted order, according to their commit system change  
number, which is to say, in the original order in which the transactions  
committed.

Moreover, the end user or the subscriber application 121 need not be  
aware of the fact that there is a join between the change table (that  
does not contain the commit system change number), and the transaction  
table (that contains the commit system change number and little else).  
The feature may be attained by providing a subscriber view 251 ,  
generated on behalf of...

13/3,K/9 (Item 3 from file: 349)

DIALOG(R)File 349:PCT FULLTEXT

1994 WIPO/Univentio. All rts. reserv.

13/16 \*\*Image available\*\*

-----VALUE-INSTANCE-CONNECTIVITY-COMPUTER-IMPLEMENTED-DATABASE-----  
BASE DE DONNEES INFORMATIQUE DE VALEURS-INSTANCES-CONNECTIVITE

Patent Applicant/Assignee:

REQUIRED TECHNOLOGIES INC, 130 West 42nd Street, 21st Floor, New York, NY  
10036, US, US (Residence), US (Nationality)

Inventor(s):

TARIN Stephen A, Suite 6G, 20 East 74th Street, New York, NY 10021, US,

Legal Representative:

WEILD David III (et al) (agent), Pennie & Edmonds LLP, 1155 Avenue of the  
Americas, New York, NY 10036, US,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200244952 A1 20020606 (WO 0244952)

Application: WO 2001US47678 20011203 (PCT/WO US0147678)

Priority Application: US 2000727423 20001201  
Designated States: AE AG AL AM AT AU AZ BA BB BG BR BY BZ CA CH CN CO CR CU  
CZ DE DK DM DZ EC EE ES FI GB GD GE GH GM HR HU ID IL IN IS JP KE KG KP  
KR KZ LC LK LR LS LT LU LV MA MD MG MK MN MW MX MZ NO NZ OM PH PL PT RO  
RU SD SE SG SI SK SL TJ TM TR TT TZ UA UG UZ VN YU ZA ZW  
(EP) AT BE CH CY DE DK ES FI FR GB GR IE IT LU MC NL PT SE TR  
(OA) BF BJ CF CG CI CM GA GN GQ GW ML MR NE SN TD TG  
(AP) GH GM KE LS MW MZ SD SL SZ TZ UG ZM ZW  
(EA) AM AZ BY KG KZ MD RU TJ TM  
Publication Language: English  
Filing Language: English  
Fulltext Word Count: 28571

Fulltext Availability:  
Detailed Description

#### Detailed Description

... for the Parity column, the "FRENCH" value corresponding to the record is the value in the corresponding cell of the FRENCH column in the value table. In the example, the entry in row 5 of the Parity column of the instance table is associated with the record being reconstructed. Thus, the "French" value is found in row 5 of the "French" column of the value table, whose value is "Trois".

Alternatively, an unsorted column may be included in the data structures of the present invention by using the identity permutation as the permutation for that column (i.e., the value table for that column will not be reordered in any way).

1 5

#### Column Merge Compression

In accordance with a farther embodiment of the invention, separate value table columns can be merged into a single column, referred to herein as a "union column," with separate displacement list columns for each of the original columns. This has the potential advantages of having a smaller value table, pre-joined data expediting join operations and improved update speed. A value not present in a particular original column is indicated in the displacement table column by a null range for that value. For example (assuming a "first row number" format displacement table), if the original column did not have the value at rcwY of the merged column, the displacement table for that column would have the same value at row Y and row T+1' (that is Displacement - Table (r+1,c)Displacement -Table(r,c)=0). If Y is the last row in the column, its value is set to a number greater...

13/3,K/10 (Item 4 from file: 349)  
ALOG(R)File 349:PCT FULLTEXT  
WIPO/Univentio. All rts. reserv.

00411827 \*\*Image available\*\*

#### AUTOMATICALLY DETERMINING A RESPONSE TO AN INQUIRY USING STRUCTURED INFORMATION

PROCEDE PERMETTANT DE DETERMINER AUTOMATIQUEMENT UNE REPONSE A UNE INTERROGATION A L'AIDE D'INFORMATIONS STRUCTUREES

Patent Applicant/Assignee:

FACT CITY INC., 245 Winter Street, Suite 220, Waltham, MA 02451, US, US  
(Residence), US (Nationality), (For all designated states except: US)

Patent Applicant/Inventor:

TRAUB Kenneth Richard, 6 Cozy Street, Watertown, MA 02472, US, US  
(Residence), US (Nationality), (Designated only for: US)

BECKERLE Michael James, 62 Beaufort Avenue, Needham, MA 02492, US, US  
(Residence), US (Nationality), (Designated only for: US)

AFERING Eric Adam, 16 Kingman Road, Newton, MA 02461, US, US (Residence),  
US (Nationality), (Designated only for: US)

Legal Representative:

FACT CITY INC (commercial rep.), 245 Winter Street, Suite 220, Waltham,  
MA 02451-8735, US,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200165412 A2-A3 20010907 (WO 0165412)  
Application: WO 2001US6342 20010228 (PCT/WO US01006342)  
Priority Application: US 2000186083 20000229; US 2000562465 20000501  
Parent Application/Grant:  
Related by Continuation to: US 2000562465 20000505 (CON); US 2000186083  
20000229 (CON)  
Designated States: AU CA US  
(EP) AT BE CH CY DE DK ES FI FR GB GR IE IT LU MC NL PT SE TR  
Publication Language: English  
Filing Language: English  
Fulltext Word Count: 15249

Fulltext Availability:  
Detailed Description

Detailed Description  
... 0, "Country" 312, and "Language" 314 columns.

All of the above Figure 2 and Figure 3 elements discussed are real elements.

They exist inside relational **tables** managed by the system. The system also provides for other types of identifiers, namely virtual elements. Virtual elements are identifiers that are not real elements...

...for example, "population", "age", and "religion". The second category of virtual element, complex virtual elements, represent information that is computed by the system from table join operations or other types of arithmetic processing.

..

For example, the "cast list" 914 associated with a movie, such as the movie titled "Rain Man", can be queried like an attribute of a particular movie. No such attribute exists as a real element inside any of the movie related tables of the illustrative system I 00. The system I 00, can perform a **table join** of the Movie Information **Table** 202 and the Movie Personnel Table I I 00 to match a **movie** identifiers in the columns 204 and I 1 04 from each table.

Figure 1 1 is a block diagram depicting a Movie Personnel Information Table I I 00. The role identifier column 1 102, lists role identifiers that identify roles associated with various movies as indicated by the movie identifiers residing in the same row and I 0 listed in the movie identifier column 1 104. The movie identifier column II 04 lists identifiers to various movies. The person identifier column 1 106 lists person identifiers of persons associated with movie roles. The role type column 1 108 lists the types of roles, such...

13/3,K/11 (Item 5 from file: 349)  
DIALOG(R)File 349:PCT FULLTEXT  
(c) 2004 WIPO/Univentio. All rts. reserv.

00557625 \*\*Image available\*\*  
ANALYTIC LOGICAL DATA MODEL  
MODELE ANALYTIQUE DE DONNEES LOGIQUES

Applicant/Assignee:

MILLER CORPORATION,  
MILLER Timothy Edward,  
TATE Brian Don,  
ROLLINS Anthony Lowell,

Inventor(s):

MILLER Timothy Edward,  
TATE Brian Don,  
ROLLINS Anthony Lowell,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200020998 A1 20000413 (WO 0020998)  
Application: WO 99US23019 19991001 (PCT/WO US9923019)  
Priority Application: US 98102831 19981002

Designated States: AE AL AM AT AU AZ BA BB BG BR BY CA CH CN CR CU CZ DE DK  
DM EE ES FI GB GD GE GH GM HR HU ID IL IN IS JP KE KG KP KR KZ LC LK LR  
LS LT LU LV MD MG MK MN MW MX NO NZ PL PT RO RU SD SE SG SI SK SL TJ TM  
TR TT TZ UA UG US UZ VN YU ZA ZW GH GM KE LS MW SD SL SZ TZ UG ZW AM AZ  
BY KG KZ MD RU TJ TM AT BE CH CY DE DK ES FI FR GB GR IE IT LU MC NL PT  
SE BF BJ CF CG CI CM GA GN GW ML MR NE SN TD TG

Publication Language: English  
Fulltext Word Count: 11217

Fulltext Availability:  
Detailed Description

Detailed Description

... the function is computed using all preceding rows. The first two rows  
in the table or group however will have the NULL value.

#### TRANINT Results Table

This table supports the TRANINT function that derives counts, amounts,  
percentage means, and intensities from a transaction summary file. This  
function takes as parameters a...

... calculate the average periodic transaction counts and transaction  
amounts by customer and transaction type (such as debit or credit) from a  
transaction summary table.

#### Data Reorganization - Results Tables and Column Definitions

The following describes the results tables and column definitions for  
the Data Reorganization functions.

#### JOIN Results Table

This table supports the JOIN function that joins tables together  
into a

combined result table. This function takes as parameters a list of  
tables, keys, and column lists to combine new derived variables into an  
analytic data set, prior to building a matrix or...

... be requested, which returns rows for all key column values found in the  
first table specified, and fills in any missing values from the other  
tables with null values.

#### DENORM Results Table

This table supports the DENORM function that selects or creates a new  
denormalized table. This function takes as parameters a table name, the  
name of key column...

13/3,K/12 (Item 6 from file: 349)

DIALOG(R)File 349:PCT FULLTEXT

(c) 2004 WIPO/Univentio. All rts. reserv.

... 9952 \*\*Image available\*\*

VALUE-INSTANCE-CONNECTIVITY COMPUTER-IMPLEMENTED DATABASE

BASE DE DONNEES INFORMATIQUE VALEUR-INSTANCE-CONNECTIVITE

Patent Applicant/Assignee:

REQUIRED TECHNOLOGIES INC,

Inventor(s):

-----  
TARIN Stephen A,

Patent and Priority Information (Country, Number, Date):

Patent: WO 200003335 A1 20000120 (WO 0003335)

Application: WO 99US15431 19990708 (PCT/WO US9915431)

Priority Application: US 98112078 19980708

Designated States: AE AL AM AT AU AZ BA BB BG BR BY CA CH CN CU CZ DE DK EE

ES FI GB GD GE GH GM HR HU ID IL IN IS JP KE KG KP KR KZ LC LK LR LS LT

LU LV MD MG MK MN MW MX NO NZ PL PT RO RU SD SE SG SI SK SL TJ TM TR TT

UA UG UZ VN YU ZA ZW GH GM KE LS MW SD SL SZ UG ZW AM AZ BY KG KZ MD RU

TJ TM AT BE CH CY DE DK ES FI FR GB GR IE IT LU MC NL PT SE BF BJ CF CG

CI CM GA GN GW ML MR NE SN TD TG

Publication Language: English

Fulltext Word Count: 26267

Fulltext Availability:  
Detailed Description

#### Detailed Description

... for the Parity column, the "FRENCH" value corresponding to the record is the value in the corresponding cell of the FRENCH column in the value table. In the example, the entry in row 5 of the Parity column of the instance table is associated with the record being reconstructed. Thus, the "French" value is found in row 5 of the "French" column of the value table, whose value is "Trois".

Alternatively, an unsorted column may be included in the data structures of the present invention by using the identity permutation as the permutation for that column (i.e., the value table for that column will not be reordered in any way).

#### Column Merge Compression

In accordance with a further embodiment of the invention, separate value table columns can be merged into a single column, referred to herein as a "union column," with separate displacement list columns for each of the original columns. This has the potential advantages of having a smaller value table, pre-joined data expediting join operations and improved update speed. A value not present in a particular original column is indicated in the displacement table column by a null range for that value. For example (assuming a "first row number" format displacement table), if the original column did not have the value at row Irl of the merged column, the displacement table for that column would have the same value at row Irl and row Ir+1 (that is Displacement-Table(r+1,c)-Displacement-Table(r,c)=0). If Irl is the last row in the column, its value is set to a number - 32...

13/3,K/13 (Item 7 from file: 349)  
DIALOG(R)File 349:PCT FULLTEXT  
(c) 2004 WIPO/Univentio. All rts. reserv.

00376923

#### STRUCTURED FOCUSED HYPERTEXT DATA STRUCTURE

#### STRUCTURE DE DONNEES HYPERTEXTE ARTICULEE SUR LA STRUCTURATION

Patent Applicant/Assignee:

HYPERMED LTD,  
OREN Avraham,  
OLCHA Lev,  
KOWALSKI Nahum,  
MARGULYAN Rita,

Inventor(s):

OREN Avraham,  
OLCHA Lev,  
KOWALSKI Nahum,  
MARGULYAN Rita,

Patent and Priority Information (Country, Number, Date):

Patent: WO 9717666 A2 19970515  
Application: WO 9611131 19961023 (PCT/WO IL9600131)  
Priority Application: US 95551929 19951023

Designated States: AL AM AT AU AZ BB BG BR BY CA CH CN CZ DE DK EE ES FI GB  
GE HU IS JP KE KG KP KR KZ LK LR LS LT LU LV MD MG MK MN MW MX NO NZ PL  
PT RO RU SD SE SG SI SK TJ TM TR TT UA UG US UZ VN KE LS MW SD SZ UG AM  
AZ BY KG KZ MD RU TJ TM AT BE CH DE DK ES FI FR GB GR IE IT LU MC NL PT  
SE BF BJ CF CG CI CM GA GN ML MR NE SN TD TG

Publication Language: English

```

TableDef,
IfIsNumberOrEnglishLetter = tblFromn As Table, tblTo As Table,
IfIsEnglishLetter(Character) Or FieldNotForCompare( As String) As
IfIsNumber(Character) Integer
End Function
Dim FieldCount As Integer
Function IsInArrayStr (ByVal aStr As String, Dim FieldNumber As Integer
String, aArray( As String) As...

```

13/3,K/14 (Item 8 from file: 349)  
DIALOG(R)File 349:PCT FULLTEXT  
(c) 2004 WIPO/Univentio. All rts. reserv..

0000256 \*\*Image available\*\*  
IMPROVED METHOD AND APPARATUS FOR DATA ACCESS IN MULTIPROCESSOR DIGITAL  
DATA PROCESSING SYSTEMS  
PROCEDE ET APPAREIL AMELIORES D'ACCES AUX DONNEES DANS DES SYSTEMES DE  
DONNEES NUMERIQUES A PROCESSEURS MULTIPLES  
Patent Applicant/Assignee:  
KENDALL SQUARE RESEARCH CORPORATION,  
Inventor(s):  
REINER David,  
MILLER Jeffrey M,  
WHEAT David C,  
Patent and Priority Information (Country, Number, Date):  
Patent: WO 9521407 A2 19950810  
Application: WO 95US1356 19950131 (PCT/WO US9501356)  
Priority Application: US 94189497 19940131  
Registered States: CA JP AT BE CH DE DK ES FR GB GR IE IT LU MC NL PT SE  
Publication Language: English  
Fulltext Word Count: 60951

Fulltext Availability:  
Claims

#### Claim

... they are executed during an initial set-up phase of query execution.  
They fall into two general groups: DDL set-up queries to create temporary  
tables or indexes; and DML set-up queries, which could be used in  
multi-stage execution strategies to populate temporary tables with  
intermediate results. Potentially, a redundant sorting of the non-driving  
table in the join by each parallel subquery, either by pre-sorting or  
by preindexing the non-driving table. If pre-sorting is used, only those  
rows which satisfy singletable predicates are inserted in a temporary  
table, which is indexed on the join columns, and  
the temporary table replaces the original table in t

he FROM clauses of the parallel  
subqueries. Ifpre-indexing' is used, the entire table must be indexed  
on the join columns. Either way, the resulting table can now be  
used as the inner table in a nested loops join. Any set-up queries  
which are generated as part of the transformation of a given query must  
be executed to completion before proceeding with execution...

...the remaining query types, and could conceptually be performed in - - - - -  
parallel with it- - - - -

#### Clean-up Queries

For each set-up query which creates a temporary table or index, a  
corresponding clean-up query is required to dispose of that temporary  
object. Clean-up queries are generated at the same time set...

13/3,K/15 (Item 9 from file: 349)  
DIALOG(R)File 349:PCT FULLTEXT  
(c) 2004 WIPO/Univentio. All rts. reserv..

00291246 \*\*Image available\*\*

METHOD AND APPARATUS FOR PARALLEL PROCESSING IN A DATABASE SYSTEM  
PROCEDE ET APPAREIL DE TRAITEMENT EN PARALLELE DANS UN SYSTEME DE BASE DE  
DONNEES

Patent Applicant/Assignee:  
ORACLE CORPORATION,

Inventor(s):  
HALLMARK Gary,  
LEARY Daniel,

Patent and Priority Information (Country, Number, Date):

Patent: WO 9509395 A1 19950406

Application: WO 94US10092 19940909 (PCT/WO US9410092)

Priority Application: US 93585 19930927

Designated States: AM AT AU BB BG BR BY CA CH CN CZ DE DK ES FI GB GE HU JP  
KE KG KP KR KZ LK LR LT LU LV MD MG MN MW NL NO NZ PL PT RO RU SD SE SI  
SK TJ TT UA UZ VN KE MW SD AT BE CH DE DK ES FR GB GR IE IT LU MC NL PT  
SE BF BJ CF CG CI CM GA GN ML MR NE SN TD TG

Publication Language: English

Fulltext Word Count: 16826

Fulltext Availability:

Detailed Description

Detailed Description

... The present invention, the smallest constituent parts are row sources.  
A row source is an objectoriented mechanism for manipulating rows of data  
in a relational database  
system (RDBMS). A row source is implemented as an iterator. Every row  
source has class methods associated with it (e.g., open, fetch next and  
dose).

Examples of row sources include: count, filter, join , sort , union,  
and table scan.

Other row sources can be used without exceeding the scope of the  
present  
invention.

As a result of the compilation process, a plan for the execution of a  
query  
is generated. An execution plan is a plan for the execution...

---

File 8: Ei Compendex(R) 1970-2004/Apr W2  
(c) 2004 Elsevier Eng. Info. Inc.  
File 35: Dissertation Abs Online 1861-2004/Mar  
(c) 2004 ProQuest Info&Learning  
File 65: Inside Conferences 1993-2004/Apr W4  
(c) 2004 BLDSC all rts. reserv.  
File 2: INSPEC 1969-2004/Apr W3  
(c) 2004 Institution of Electrical Engineers  
File 94: JICST-EPlus 1985-2004/Apr W2  
(c) 2004 Japan Science and Tech Corp(JST)  
File 6: NTIS 1964-2004/Apr W4  
(c) 2004 NTIS, Intl Cpyrghrt All Rights Res  
File 144: Pascal 1973-2004/Apr W3  
(c) 2004 INIST/CNRS  
File 434: SciSearch(R) Cited Ref Sci 1974-1989/Dec  
(c) 1998 Inst for Sci Info  
File 34: SciSearch(R) Cited Ref Sci 1990-2004/Apr W3  
(c) 2004 Inst for Sci Info  
File 99: Wilson Appl. Sci & Tech Abs 1983-2004/Mar  
(c) 2004 The HW Wilson Co.  
File 583: Gale Group Globalbase(TM) 1986-2002/Dec 13  
(c) 2002 The Gale Group  
File 266: FEDRIP 2004/Feb  
Comp & dist by NTIS, Intl Copyright All Rights Res  
File 95: TEME-Technology & Management 1989-2004/Apr W1  
(c) 2004 FIZ TECHNIK  
File 62: SPIN(R) 1975-2004/Feb W5  
(c) 2004 American Institute of Physics  
File 239: Mathsci 1940-2004/Jun  
(c) 2004 American Mathematical Society

Set	Items	Description
S1	4083	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (7N) (ROW? ? OR TUPLE? ? OR COLUMN? ? OR FIELD? ? OR CELL? ? OR TABLE? ?)
S2	365	S1(7N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIME OR AS OR WHILE)
S3	1221004	DATABASE? ? OR DATA()BASE? ? OR TABLE? ?
S4	4190	(JOIN OR JOINS OR JOINED OR JOINING OR EQUIJOIN???) (20N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC OR SAME()TIME)
	824	S1(25N) (SIMULTANEOUS? OR CONCURREN? OR COINCIDENT? OR PARALLEL OR SYNCHRONIZ? OR SYNCHRONIS??? OR SYNCHRONOUS? OR SYNC - OR SAME()TIME OR AS OR WHILE)
S6	95207	REDISTRIBUT? OR RE()DISTRIBUT?
S7	50	S4:S5 AND S6
S8	36	S3 AND S7
S9	21	RD (unique items)
S10	17	S9 NOT PY=2002:2004
S11	83	S2 AND S3
S12	18143573	ROW? ? OR TUPLE? ? OR COLUMN? ? OR FIELD? ? OR CELL? ? OR - DATA OR INFORMATION
S13	590701	S12(5N) (REDISTRIBUT? OR DISTRIBUT? OR SHUFFL? OR RESHUFFL? OR SHIFT??? OR RESHIFT??? OR MOV??? OR MIX??? OR REORDER??? OR SORT??? OR RESORT??? OR REARRANG? OR REORGANI? OR REGROUP??? OR RE() (ARRANG? OR ORGANI? OR GROUP??? OR ORDER?))
S14	15	S11 AND S13
S15	10	RD (unique items)
S16	29	S5 AND S3 AND S13
S17	19	RD (unique items)
S18	19	S15 OR S17
S19	11	S18 NOT (S10 OR PY=2002:2004)

10/5/1 (Item 1 from file: 8)  
DIALOG(R) File 8: Ei Compendex(R)  
(c) 2004 Elsevier Eng. Info. Inc. All rts. reserv.

04623685 E.I. No: EIP97023519632

Title: Partition based spatial-merge join  
Author: Patel, Jignesh M.; DeWitt, David J.  
Corporate Source: Univ of Wisconsin, Madison, WI, USA  
Conference Title: Proceedings of the 1996 ACM SIGMOD International  
Conference on Management of Data  
Conference Location: Montreal, Can Conference Date: 19960604-19960606  
Sponsor: ACM SIGMOD  
E.I. Conference No.: 45963  
Source: SIGMOD Record (ACM Special Interest Group on Management of Data)  
June 1996.. p 259-270  
Publication Year: 1996  
CODEN: SRECD8  
Language: English  
Document Type: CA; (Conference Article) Treatment: G; (General Review);  
T; (Theoretical)

Journal Announcement: 9704W1

Abstract: This paper describes PBSM (Partition Based Spatial-Merge), a new algorithm for performing spatial join operation. This algorithm is especially effective when neither of the inputs to the join have an index on the joining attribute. Such a situation could arise if both inputs to the join are intermediate results in a complex query, or in a parallel environment where the inputs must be dynamically redistributed. The PBSM algorithm partitions the inputs into manageable chunks, and joins them using a computational geometry based plane-sweeping technique. This paper also presents a performance study comparing the traditional indexed nested loops join algorithm, a spatial join algorithm based on joining spatial indices, and the PBSM algorithm. These comparisons are based on complete implementations of these algorithms in Paradise, a database system for handling GIS applications. Using real data sets, the performance study examines the behavior of these spatial join algorithms in a variety of situations, including the cases when both, one, or none of the inputs to the join have a suitable index. The study also examines the effect of clustering the join inputs on the performance of these join algorithms. The performance comparisons demonstrate the feasibility, and applicability of the PBSM join algorithm. (Author abstract) Refs.

Descriptors: Merging; Algorithms; Query languages; Computational geometry; Data acquisition; Relational database systems; Geographic information systems

Identifiers: Partition based spatial merge (PBSM) join; Plane sweeping techniques; Nested loops join algorithm; Spatial indices; Database system paradise

Classification Codes:

723.1.1 (Computer Programming Languages)  
723.1 (Computer Programming); 723.2 (Data Processing); 723.3 (Database Systems)  
723 (Computer Software); 921 (Applied Mathematics)  
72 (COMPUTERS & DATA PROCESSING); 92 (ENGINEERING MATHEMATICS)

10/5/2 (Item 2 from file: 8)  
DIALOG(R) File 8: Ei Compendex(R)  
--- (c) 2004 Elsevier Eng. Info. Inc. All rts. reserv.

04444667 E.I. No: EIP96073246434

Title: Parallel transitive closure computation in relational databases  
Author: Zhou, Xiaofang; Zhang, Yanchun; Orlowska, Maria E.  
Corporate Source: CSIRO Div of Information Technology, Canberra, Aust  
Source: Information Sciences v 92 n 1-4 Jul 1996. p 109-135  
Publication Year: 1996  
CODEN: ISIJBC ISSN: 0020-0255  
Language: English  
Document Type: JA; (Journal Article) Treatment: A; (Applications); T; (Theoretical)

Journal Announcement: 9609W2

Abstract: The transitive closure operation is an important extension to relational algebra. Because of its high computation cost, it is of great interest to design efficient parallel algorithms for computing the transitive closure in relational database systems. In this paper, we present a new algorithm to compute transitive closures on SIMD meshes based on relational algebra operations. Double-hash distribution is used to avoid rehashing new tuples for the next join phase. There presently exists no extra step for the redistribution of these tuples. Possible redundant computation between different join phases has been prevented without using global operations. As only regular linear communication occurs on the mesh, and the workload is fully distributed, a speedup of  $O(n)$  (multiplied by  $n$ ) has been achieved, where  $n$  multiplied by  $n$  is the size of mesh. Therefore, this algorithm is an optimal parallel version of the transitive closure algorithms based on relational algebra operations on SIMD meshes. (Author abstract) 28 Refs.

Descriptors: Parallel algorithms; Relational database systems; Algebra; Parallel processing systems; Computational methods; Optimization; Distributed computer systems

Identifiers: Parallel transitive closure computation; Relational algebra; Double hash distribution; Global operations; Regular linear communication; Single input multiple data

Classification Codes:

723.1 (Computer Programming); 723.3 (Database Systems); 921.1 (Algebra); 722.4 (Digital Computers & Systems); 721.1 (Computer Theory, Includes Formal Logic, Automata Theory, Switching Theory, Programming Theory); 921.5 (Optimization Techniques)

723 (Computer Software); 921 (Applied Mathematics); 722 (Computer Hardware); 721 (Computer Circuits & Logic Elements)

72 (COMPUTERS & DATA PROCESSING); 92 (ENGINEERING MATHEMATICS)

10/5/3 (Item 3 from file: 8)

DIALOG(R) File 8: Ei Compendex(R)

(c) 2004 Elsevier Eng. Info. Inc. All rts. reserv.

9609W2 E.I. No: EIP93040764372

Title: Join and data redistribution algorithms for hypercubes

Author: Baru, Chaitanya K.; Padmanabhan, Sriram

Source: IEEE Transactions on Knowledge and Data Engineering v 5 n 1 Feb 1993. p 161-168

Publication Year: 1993

CODEN: ITKEEH ISSN: 1041-4347

Language: English

Document Type: JA; (Journal Article) Treatment: T; (Theoretical)

Journal Announcement: 9307W1

Abstract: An important aspect of database processing in parallel computer systems is the use of data parallel algorithms. This paper presents several parallel algorithms for the relational database join operation in a hypercube multicomputer system. The join algorithms are classified as cycling or global partitioning based on the tuple distribution method employed. The various algorithms are compared under a common framework, using time complexity analysis as well as an implementation on a 64 node NCUBE hypercube system. In general, the global partitioning algorithms demonstrate better speedup. However, the cycling algorithm can perform better than the global algorithms in specific situations, viz, when the difference in input relation cardinalities is large and the hypercube dimension is small. We also study the usefulness of the data redistribution operation in improving the performance of the join algorithms, in the presence of uneven data partitions. Our results indicate that redistribution significantly decreases the join algorithm execution times for unbalanced partitions. (Author abstract) 21 Refs.

Descriptors: Algorithms; Parallel processing systems; Multiprocessing; Data processing

Identifiers: Data redistribution algorithms; Hypercubes; Cycling algorithms; Time complexity; Global algorithms; Join algorithms

Classification Codes:

723.1 (Computer Programming); 723.3 (Database Systems); 722.4 (Digital

Computers & Systems); 723.2 (Data Processing)  
723 (Computer Software); 722 (Computer Hardware)  
72 (COMPUTERS & DATA PROCESSING)

10/5/4 (Item 4 from file: 8)  
DIALOG(R)File 8:EI Compendex(R)  
(c) 2004 Elsevier Eng. Info. Inc. All rts. reserv.

02817938 E.I. Monthly No: EI8911114257

Title: Database operations in a cube-connected multicomputer system.

Author: Baru, Chaitanya K.; Frieder, Ophir

Corporate Source: Univ of Michigan, Ann Arbor, MI, USA

Source: IEEE Transactions on Computers v 38 n 6 Jun 1989 p 920-927

Publication Year: 1989

CODEN: ITCOB4 ISSN: 0018-9340

Language: English

Document Type: JA; (Journal Article) Treatment: T; (Theoretical)

Journal Announcement: 8911

Abstract: Distributed memory architectures, specifically hypercubes, for parallel database processing are treated. The cube interconnects support efficient data combination for the various database operations, and nonuniform data distributions are handled by dynamically redistributing data utilizing these interconnections. Selection and scalar aggregation operations are easily supported. An algorithm for the join operation is discussed in some detail. The cube is compared with another multicomputer database machine, SM3, and the performance of the join operation in these systems is described. The join performance in a cube is comparable to that in SM3 even when the cube is assumed to have a nonuniform data distribution. 15 refs.

Descriptors: COMPUTER SYSTEMS, DIGITAL--\*Parallel Processing; DATABASE SYSTEMS--Distributed; COMPUTER PROGRAMMING--Algorithms

Identifiers: DISTRIBUTED MEMORY ARCHITECTURES; PARALLEL DATABASE PROCESSING; DATABASE OPERATIONS; CUBE CONNECTED MULTIPROCESSORS; JOIN ALGORITHMS

Classification Codes:

722 (Computer Hardware); 723 (Computer Software)

72 (COMPUTERS & DATA PROCESSING)

10/5/5 (Item 5 from file: 8)  
DIALOG(R)File 8:EI Compendex(R)  
(c) 2004 Elsevier Eng. Info. Inc. All rts. reserv.

02264684 E.I. Monthly No: EIM8708-051934

Title: IMPLEMENTING RELATIONAL DATABASE OPERATIONS IN A CUBE-CONNECTED MULTICOMPUTER SYSTEM.

Author: Baru, Chaitanya K.; Frieder, Ophir

Corporate Source: Univ of Michigan, Ann Arbor, MI, USA

Conference Title: Proceedings - Third International Conference on Data Engineering.

Conference Location: Los Angeles, CA, USA Conference Date: 19870203

Sponsor: IEEE Computer Soc, Los Alamitos, CA, USA

E.I. Conference No.: 09779

Source: Publ by IEEE, New York, NY, USA. Available from IEEE Service Cent

--- (Cat\_n 87CH2407-5), Piscataway, NJ, USA-p-36-43 ---

Publication Year: 1987

ISBN: 0-8186-0762-9

Language: English

Document Type: PA; (Conference Paper)

Journal Announcement: 8708

Abstract: Strategies for performing database operation in a cube-connected multicomputer system with parallel I/O are presented. The cube interconnection subsumes many other structures such as the tree, ring, etc. This property is utilized to efficiently support database operations such as select, aggregate, join, and project. The strategies presented are unique in that they account for the nonuniform distribution of data across parallel paths by incorporating data redistribution steps as

part of the overall algorithm. The two main data redistribution operations used are tuple balancing and merging. 17 refs.

Descriptors: DATABASE SYSTEMS--\*Relational; COMPUTER SYSTEMS, DIGITAL--Multiprocessing

Identifiers: MULTICOMPUTER SYSTEMS; CUBE INTERCONNECTIONS; DATA DISTRIBUTION

Classification Codes:

723 (Computer Software); 722 (Computer Hardware)

72 (COMPUTERS & DATA PROCESSING)

10/5/6 (Item 1 from file: 35)

DIALOG(R)File 35:Dissertation Abs Online

1994 ProQuest Info&Learning. All rts. reserv.

ORDER NO: AAD97-34729

# QUERY PROCESSING IN SPATIAL DATABASE SYSTEMS: DECLUSTERING AND CLUSTERING TECHNIQUES

Author: RAVADA, SIVAKUMAR

Degree: PH.D.

Year: 1997

Corporate Source/Institution: UNIVERSITY OF MINNESOTA (0130)

Source: VOLUME 58/05-B OF DISSERTATION ABSTRACTS INTERNATIONAL.

PAGE 2515. 94 PAGES

Descriptors: COMPUTER SCIENCE

Descriptor Codes: 0984

The research question in this thesis concerns how to parallelize the spatial range and join query processing in order to support a high performance spatial database application. Data partitioning for the range query operation involves declustering of spatial data, while data partitioning for the spatial join involves clustering of spatial data. If the static partitioning methods fail to equally distribute the load among different processors, the load-balance may be improved by redistributing parts of the data to idle processors using Dynamic Load-Balancing (DLB) techniques.

In this thesis, we provide a framework for declustering collections of extended spatial objects by identifying the following key issues: (i) the work-load metric, (ii) the spatial-extent of the work-load, (iii) the distribution of the work-load over the spatial-extent, and (iv) the declustering method. We identify and experimentally evaluate alternatives to these issues.

In addition, we also provide a framework for dynamically balancing the load between different processors. We experimentally evaluate the proposed declustering and load-balancing methods on a distributed memory MIMD machine (Cray T3D) and shared-memory machine (SGI Challenge). Experimental results show that the spatial-extent and the work-load metric are important issues in developing a declustering method. Experiments also show that the replication of data is usually needed to facilitate dynamic load-balancing, as the cost of local processing is often less than the cost of data transfer for extended spatial objects. In addition, we also show that the effectiveness of dynamic load-balancing techniques can be improved by using declustering methods to determine the subsets of spatial objects to be transferred during run-time.

A spatial join is often performed in two steps: a filter step and a refinement step. In this thesis, we focus on the refinement step of the spatial join. The refinement step of the spatial join takes as input a sequence of pairs of tuples and checks each tuple to see if the join predicate is satisfied for that tuple. This is similar to the join index processing done in traditional relational databases. We develop min-cut graph partitioning based methods for join processing using a join index. We use min-cut graph partitioning as a new heuristic for solving the page access sequence problem for fixed size buffer in sequential systems. We show that the number of page accesses needed to compute a join using join index in a fixed buffer environment is bounded by the sum of sizes of the base relations and the size of the cut-set of the page connectivity graph. Since the min-cut graph partitioning aims to minimize the size of the cut-set, this proposed heuristic is a direct method. Experiments with

benchmark data sets show that the graph-partitioning based heuristic outperforms the existing methods, particularly when join selectivity is high and buffer space is small. (Abstract shortened by UMI.)

10/5/7 (Item 1 from file: 2)

DIALOG(R)File 2:INSPEC

© 2004 Institution of Electrical Engineers. All rts. reserv.

INSPEC Abstract Number: C2002-10-6160B-020

Title: AMOS-SDDS: a scalable distributed data manager for Windows multicomputers

Author(s): Ndiaye, Y.; Dilne, A.W.; Litwin, W.; Risch, T.

Author Affiliation: CERIA, Univ. Paris IX Dauphine, France

Conference Title: Proceedings of the ISCA 14th International Conference Parallel and Distributed Computing Systems p.523-9

Editor(s): Sha, E.

Publisher: Int. Soc. Comput. & their Applications - ISCA, Cary, NC, USA

Publication Date: 2001 Country of Publication: USA viii+529 pp.

ISBN: 1 880843 39 0 Material Identity Number: XX-2002-01770

Conference Title: Proceedings of the ISCA 14th International Conference Parallel and Distributed Computing Systems

Conference Sponsor: Int. Soc. Comput. & their Applications - ISCA

Conference Date: 8-10 Aug. 2001 Conference Location: Richardson, TX, USA

Language: English Document Type: Conference Paper (PA)

Treatment: Practical (P)

Abstract: Known parallel DBMS offer at present only static partitioning schemes. Adding a storage node is a cumbersome operation that typically requires the manual data redistribution. We present an architecture termed AMOS-SDDS for a share-nothing multicomputer. We have coupled a high-performance main-memory DBMS AMOS-II and a manager of Scalable Distributed Data Structures (SDDS) into a scalable distributed system SDDS provides the scalable data partitioning in distributed RAM, supporting parallel scans with function shipping. AMOS-SDDS couples both systems using the AMOS-II foreign function interface. Its scalability abolishes the cumbersome storage limits of a single site RAM DBMS technology. Its distributed RAM query processing and scalable data partitioning is an improvement over the current parallel DBMSs technology. We validate AMOS-SDDS architecture by experiments with distributed nested loop join queries over a file scaling up to 300,000 tuples. It includes performance study of speed-up and scale-up characteristics. The results encourage the use of SDDS for high-performance database systems. (19 Refs)

Subfile: C

Descriptors: data structures; distributed databases ; multiprocessing systems; query processing

Identifiers: parallel DBMS; share-nothing multicomputer; distributed database ; query processing; data partitioning; distributed data structures ; RAM database systems; scalability

Class Codes: C6160B (Distributed databases); C6120 (File organisation)

Copyright 2002, IEE

10/5/8 (Item 2 from file: 2)

DIALOG(R)File 2:INSPEC

© 2004 Institution of Electrical Engineers.-All-rts.-reserv.

7012626 INSPEC Abstract Number: C2001-09-4250-035

Title: A skew-insensitive algorithm for join and multi-join operations on shared nothing machines

Author(s): Bamha, M.; Hains, G.

Author Affiliation: LIF, Univ. d, France

Conference Title: Database and expert systems applications. 11th International Conference, DEXA 2000. Proceedings (Lecture Notes in Computer Science Vol.1873) p.644-53

Editor(s): Ibrahim, M.; Kung, J.; Revell, N.

Publisher: Springer-Verlag, Berlin, Germany

Publication Date: 2000 Country of Publication: Germany xix+1003 pp.

ISBN: 3 540 67978 2      Material Identity Number: XX-2001-01532  
Conference Title: Database and Expert Systems Applications. 11th  
International Conference, DEXA 2000. Proceedings  
Conference Date: 4-8 Sept. 2000      Conference Location: London, UK  
Language: English      Document Type: Conference Paper (PA)  
Treatment: Practical (P)

Abstract: Join is an expensive and frequently used operation whose parallelization is highly desirable. However, the effectiveness of parallel joins depends on the ability to evenly divide load among processors. Data skew can have a disastrous effect on performance. Although many skew-handling algorithms have been proposed, they remain generally inefficient in the case of multi-joins due to join product skew, costly and unnecessary redistribution and communication costs. A parallel join algorithm called fa-join was introduced in an earlier paper with deterministic and near-perfect balancing properties. Despite its advantages, fa-join is sensitive to the correlation of the attribute value distributions in both relations. We present an improved version of the algorithm called Sfa-join with a symmetric treatment of both relations. Its remarkably low join-product and attribute-value skew makes it suitable for repeated use in multi-join operations. Its performance is analyzed theoretically and experimentally, to confirm its linear speed-up and its superiority over fa-join. (16 Refs)

Subfile: C

Descriptors: parallel algorithms; parallel databases; relational algebra; resource allocation

Identifiers: skew-insensitive algorithm; multi-join operations; join operations; shared nothing machines; join parallelization; parallel joins; data skew; skew-handling algorithms; multi-joins; join product skew; communication costs; fa-join; near-perfect balancing properties; attribute value distributions; Sfa-join; symmetric treatment; attribute-value skew; linear speed-up

Class Codes: C4250 (Database theory); C6160D (Relational databases); C6160B (Distributed databases); C6110P (Parallel programming); C4240P (Parallel programming and algorithm theory)

Copyright 2001, IEE

10/5/9      (Item 3 from file: 2)  
DIALOG(R)File 2:INSPEC  
(c) 2004 Institution of Electrical Engineers. All rts. reserv.

6571972      INSPEC Abstract Number: C2000-06-6160B-002

Title: On disk allocation of intermediate query results in parallel database systems

Author(s): Martens, H.

Author Affiliation: Inst. fur Inf., Leipzig Univ., Germany

Conference Title: Euro-Par'99. Parallel Processing. 5th International Euro-Par Conference. Proceedings (Lecture Notes in Computer Science Vol.1685) p.469-76

Editor(s): Amestoy, P.; Berger, P.; Dayde, M.; Duff, I.; Frayssé, V.; Giraud, L.; Ruiz, D.

Publisher: Springer-Verlag, Berlin, Germany

Publication Date: 1999      Country of Publication: Germany      xxxii+1503 pp.

ISBN: 3 540 66443 2      Material Identity Number: XX-1999-02565

Conference Title: Proceedings of Euro-Par'99

Conference Date: 31 Aug.-3 Sept. 1999      Conference Location: Toulouse, France

Language: English      Document Type: Conference Paper (PA)

Treatment: Practical (P)

Abstract: For complex queries in parallel database systems, substantial amounts of data must be redistributed between operators executed on different processing nodes. Frequently, such intermediate results cannot be held in main memory and must be stored on disk. To limit the ensuing performance penalty, a data allocation must be found that supports parallel I/O to the greatest possible extent. In this paper, we propose declustering even self-contained units of temporary data processed in a single operation (such as individual buckets of parallel hash joins) across multiple disks. Using a suitable analytical model, we find that the improvement of

parallel I/O outweighs the penalty of increased fragmentation. (14 Refs)  
Subfile: C  
Descriptors: data warehouses; file organisation; parallel databases  
Identifiers: disk allocation; intermediate query results; parallel  
database systems; complex queries; performance penalty; data allocation;  
declustering; parallel hash joins  
Class Codes: C6160B (Distributed databases); C6120 (File organisation)  
Copyright 2000, IEE

10/5/10 (Item 4 from file: 2)  
DIALOG(R)File 2:INSPEC  
(c) 2004 Institution of Electrical Engineers. All rts. reserv.

6205550 INSPEC Abstract Number: C1999-05-6160B-004  
Title: Performance of load balancing techniques for join operations in  
shared-nothing database management systems  
Author(s): Hua, K.A.; Tavanapong, W.; Yu-Lung Lo  
Author Affiliation: Sch. of Comput. Sci., Central Florida Univ., Orlando,  
FL, USA

Journal: Journal of Parallel and Distributed Computing vol.56, no.1  
p.17-46

Publisher: Academic Press,  
Publication Date: Jan. 1999 Country of Publication: USA  
CODEN: JPD CER ISSN: 0743-7315  
SICI: 0743-7315(199901)56:1L:17:PLBT;1-K  
Material Identity Number: G544-1999-002  
U.S. Copyright Clearance Center Code: 0743-7315/99/\$30.00  
Language: English Document Type: Journal Paper (JP)  
Treatment: Applications (A); Practical (P)

Abstract: We investigate various load balancing approaches for hash-based  
join techniques popular in multicomputer-based shared-nothing database  
systems. When the tuples are not uniformly distributed among the hash  
buckets, redistribution of these buckets among the processors is  
necessary to maintain good system performance. Two recent load balancing  
techniques which rely on sampling and incremental balancing, respectively,  
have been shown to be more robust than conventional methods. The comparison  
of these two approaches, however, has not been investigated. In this study,  
we improve these two schemes and implement them along with a conventional  
method and a standard join technique which does not do load balancing on  
an nCUBE/2 parallel computer to compare their performance. Our  
experimental results indicate that the sampling technique is the better  
approach. To further evaluate the performance of these techniques under  
diverse hardware conditions, we also develop a cost model and implement a  
simulator to perform sensitivity analyses with respect to various hardware  
parameters. The simulation results show that both sampling and incremental  
techniques provide noticeable savings over conventional methods, with the  
sampling approach being more scalable in supporting very large database  
systems. (26 Refs)

Subfile: C  
Descriptors: parallel databases ; performance evaluation; resource  
allocation; very large databases  
Identifiers: load balancing; join operations; shared-nothing database  
management systems; hash-based join techniques; nCUBE/2 parallel computer;  
cost model; very large database systems  
Class Codes: C6160B (Distributed databases); C5470 (Performance  
evaluation and testing); C6150J (Operating systems)  
Copyright 1999, IEE

10/5/11 (Item 5 from file: 2)  
DIALOG(R)File 2:INSPEC  
(c) 2004 Institution of Electrical Engineers. All rts. reserv.

6205553 INSPEC Abstract Number: C9809-6160-003  
Title: Multiple-weighted-tree based query optimization method for parallel  
database systems  
Author(s): Li Jian-Zhong

Author Affiliation: Dept. of Comput. Sci., Heilongjiang Univ., China  
Journal: Chinese Journal of Computers vol.21, no.5 p.401-12  
Publisher: Science Press,  
Publication Date: May 1998 Country of Publication: China  
CODEN: JIXUDT ISSN: 0254-4164  
SICI: 0254-4164(199805)21:5L:401:MWTB;1-2  
Material Identity Number: B714-98007  
Language: Chinese Document Type: Journal Paper (JP)  
Treatment: Practical (P)

Abstract: A multiple-weighted-tree based query optimization method for parallel database systems is proposed in this paper. The method consists of a multiple-weighted-tree based parallel query plan model, a cost model for parallel query plans, and a query optimizer. The parallel query plan model is the first one to model all basic relational operations, all three types of parallelism of query execution in parallel database systems, processor and memory allocation to operations, memory allocation to the buffers between operations in pipelines and data redistribution among processors. The cost model takes the waiting time of the operations in pipelining execution into consideration and is computable in bottom-up fashion. The query optimizer addresses the query optimization problem in the context of Select-Project-Join queries that are widely used in commercial DBMSs. Several heuristics determining the processor allocation to operations are derived and used in the query optimizer. The query optimizer is aware of memory resources in order to generate good-quality plans. It includes the heuristics for determining the memory allocation to operations and buffers between operations in pipelines so that the memory resource is fully exploit. In addition, multiple algorithms for implementing join operations are considered in the query optimizer. The query optimizer can make an optimal choice of join algorithm for each join operation in a query. The proposed query optimization method has been used in a prototype parallel database management system designed and implemented by the author. (11 Refs)

File: C

Descriptors: database management systems; parallel machines; query processing

Identifiers: multiple-weighted-tree based query optimization; parallel database systems; cost model; parallel query plans; query optimizer; relational operations; memory allocation; Select-Project-Join queries; multiple algorithms; join algorithm; join operation

Class Codes: C6160 (Database management systems (DBMS)); C5440 (Multiprocessing systems); C4250 (Database theory)

Copyright 1998, IEE

10/5/12 (Item 6 from file: 2)

FILE: File 2:INSPEC

Institution of Electrical Engineers. All rts. reserv.

INSPEC Abstract Number: C9610-4250-002

Title: Parallel transitive closure computation in relational databases

Author(s): Xiaofang Zhou; Yanchun Zhang; Orłowska, M.E.

Author Affiliation: Div. of Inf. Technol., CSIRO, Canberra, ACT, Australia

Journal: Information Sciences vol.92, no.1-4 p.109-35

Publisher: Elsevier,

Publication Date: July 1996 Country of Publication: USA

CODEN: ISIJBC ISSN: 0020-0255

SICI: 0020-0255(199607)92:1/4L:109:PTCC;1-9

Material Identity Number: I132-96009

U.S. Copyright Clearance Center Code: 0020-0255/96/\$15.00

Document Number: S0020-0255(96)00053-9

Language: English Document Type: Journal Paper (JP)

Treatment: Theoretical (T)

Abstract: The transitive closure operation is an important extension to relational algebra. Because of its high computational cost, it is of great interest to design efficient parallel algorithms for computing the transitive closure in relational database systems. In this paper, we present a new algorithm to compute transitive closures on SIMD meshes based

on relational algebra operations. The double-hash distribution is used to avoid rehashing new tuples for the next join phase. There presently exists no extra step for the redistribution of these tuples. Possible redundant computation between different join phases has been prevented without using global operations. As only regular linear communication occurs on the mesh and the workload is fully distributed, a speedup of  $O(n \cdot n)$  has been achieved, where  $n \cdot n$  is the size of mesh. Therefore, this algorithm is an optimal parallel version of the transitive closure algorithms based on relational algebra operations on SIMD meshes. (28 Refs)

Subfile: C

Descriptors: computational complexity; database theory; parallel algorithms; relational algebra; relational databases  
Indexing: parallel transitive closure computation; relational databases; relational algebra; efficient parallel algorithms; SIMD meshes; double-hash distribution; join phases; tuple redistribution; redundant computation; linear communication; fully distributed workload; speedup; mesh size; optimal parallel version

Class Codes: C4250 (Database theory); C6160D (Relational databases); C4240P (Parallel programming and algorithm theory); C6120 (File organisation); C4240C (Computational complexity)

Copyright 1996, IEE

10/5/13 (Item 7 from file: 2)

DIALOG(R)File 2:INSPEC

(c) 2004 Institution of Electrical Engineers. All rts. reserv.

0050854 INSPEC Abstract Number: C9510-4250-018

Title: Dynamic join product skew handling for hash-joins in shared-nothing database systems

Author(s): Harada, L.; Kitsuregawa, M.

Conference Title: Database Systems for Advanced Applications '95. Proceedings of the Fourth International Conference on Database Systems for Advanced Applications p.246-55

Editor(s): Ling, T.W.; Masunaga, Y.

Publisher: World Scientific, Singapore

Publication Date: 1995 Country of Publication: Singapore xv+468 pp.

ISBN: 981 02 2220 3

Conference Title: Proceedings of 4th International Symposium on Database Systems for Advanced Applications

Conference Date: 10-13 April 1995 Conference Location: Singapore

Language: English Document Type: Conference Paper (PA)

Treatment: Theoretical (T)

Abstract: When data is uniformly distributed, the parallel hash-based join algorithm scales up well. However, the presence of data skew can cause load imbalance among the processors, significantly deteriorating its performance. In this paper we propose a dynamic skew handling algorithm which deals with this load imbalance, by detecting and handling join product skews at run-time. The idea is to monitor the join processing at the join phase and compare the average processing rate of each partition with the rate statically predicted at the scheduling phase. If their difference is detected to be large enough to produce a significant performance degradation, the processor is considered to be overloaded and a workload compensation strategy is dynamically invoked. In this case, based on the measured average processing rate, the amount of overload caused by the unpredicted join product skew is calculated and, the amount of load to be migrated to the non-overloaded processors is determined. We propose two methods-the result redistribution and the processing task migration-to handle the load migration from the overloaded processor to the non-overloaded processors. Simulation results show that our dynamic skew handling approach can detect and handle load imbalances efficiently, so that the rebalance of load among the processors results in an almost constant join execution time under different join product skews. (11 Refs)

Subfile: C

Descriptors: database theory; distributed databases; file organisation; query processing; resource allocation; software performance  
Indexing: very large databases

Identifiers: dynamic join product skew handling; hash-joins; shared-nothing database systems; parallel hash-based join algorithm; data skew; load imbalance; run-time; join processing; average processing rate; performance degradation; workload compensation strategy; unpredicted join product skew; result redistribution; processing task migration; join execution time; join product skews

Class Codes: C4250 (Database theory); C6120 (File organisation); C6160B (Distributed databases); C6150J (Operating systems)

Copyright 1995, IEE

10/5/14 (Item 8 from file: 2)

DIALOG(R)File 2:INSPEC

(c) 2004 Institution of Electrical Engineers. All rts. reserv.

5050305 INSPEC Abstract Number: C9510-6160Z-011

Title: Analysis of dynamic load balancing strategies for parallel shared nothing database systems

Author(s): Rahm, E.; Marek, R.

Author Affiliation: Kaiserslautern Univ., Germany

Conference Title: 19th International Conference on Very Large Data Bases Proceedings p.182-93

Editor(s): Agrawal, R.; Baken, S.; Bell, D.

Publisher: Morgan Kaufmann Publishers, Palo Alto, CA, USA

Publication Date: 1993 Country of Publication: USA xx+711 pp.

Conference Title: Proceeding of 19th International Conference on Very Large Data Bases

Conference Sponsor: VLD Endowment; Irish Comput. Soc.; Trinity College Dublin; et al

Conference Date: 24-27 Aug. 1993 Conference Location: Dublin, Ireland

Language: English Document Type: Conference Paper (PA)

Treatment: Practical (P)

Abstract: Parallel database systems have to support both inter-transaction as well as intra-transaction parallelism. Inter-transaction parallelism (multi-user mode) is required to achieve high throughput, in particular for OLTP transactions, and sufficient cost-effectiveness. Intra-transaction parallelism is a prerequisite for reducing the response time of complex and data-intensive transactions (queries). In order to achieve both goals dynamic strategies for load balancing and scheduling are necessary which take the current system state into account for allocating transactions and subqueries to processors and for determining the degree of intra-transaction parallelism. We study the load balancing problem for parallel join processing in shared nothing database systems. In these systems, join processing is typically based on a dynamic redistribution of relations to join processors thus making dynamic load balancing strategies feasible. In particular, we study the performance of dynamic load balancing strategies for determining the number of join processors and for selection of the join processors. In contrast to previous studies on parallel join processing, we present a multi-user performance analysis for both homogeneous and heterogeneous/mixed workloads as well as for different database allocations. (34 Refs)

Subfile: C

Descriptors: multiprocessor interconnection networks; parallel processing; processor scheduling; query processing; resource allocation; transaction processing; very large databases

Identifiers: dynamic load balancing strategies; parallel shared nothing database systems; inter-transaction parallelism; intra-transaction parallelism; cost-effectiveness; response time; data-intensive transactions; complex transactions; queries; scheduling; current system state; transaction allocation; subquery allocation; processors; parallel join processing; dynamic relation redistribution; multi-user performance analysis; homogeneous workloads; heterogeneous workloads; database allocations; mixed workloads

Class Codes: C6160Z (Other DBMS); C6150N (Distributed systems software); C6130 (Data handling techniques)

Copyright 1995, IEE

10/5/15 (Item 1 from file: 94)

DIALOG(R)File 94:JICST-EPlus

(c)2004 Japan Science and Tech Corp(JST). All rts. reserv.

02615023 JICST ACCESSION NUMBER: 96A0149622 FILE SEGMENT: JICST-E  
Dynamic Load Balancing for Right-Deep Pipelined Hash Multi- Joins for  
Shared Nothing Parallel Database Servers.

DAVIS S (1); KITSUREGAWA M (1)

(1) Univ. Tokyo, Tokyo

Denshi Joho Tsushin Gakkai Gijutsu Kenkyu Hokoku(IEIC Technical Report  
(Institute of Electronics, Information and Communication Engineers),  
1995, VOL.95,NO.410(DE95 76-84), PAGE.45-50, FIG.2, TBL.1, REF.2

JOURNAL NUMBER: S0532BBG

UNIVERSAL DECIMAL CLASSIFICATION: 681.3:061.68 681.32

LANGUAGE: English COUNTRY OF PUBLICATION: Japan

DOCUMENT TYPE: Journal

ARTICLE TYPE: Original paper

MEDIA TYPE: Printed Publication

ABSTRACT: This paper presents a dynamic load balancing algorithm for  
right-deep pipelined hash multi-joins executed on a shared nothing  
architecture. The algorithm makes use of a centralized processor called  
the foreman to gather statistics from the processors participating in  
the join. With these statistics, the foreman is able to determine how  
the build relation tuples should be redistributed to equalize the  
number of result tuples remaining to be generated by each of the  
processors. The unit of load migration is a hash line, the set of build  
tuples mapped to the same hash entry in a pipeline stage's hash table  
, and load balancing is performed separately for each stage of the  
pipeline. (author abst.)

DESCRIPTORS: DBMS; parallel processing; scheduling; hashing; parallel  
computer; algorithm; pipeline processing; computer simulation;  
distributed processing; load sharing

BROADER DESCRIPTORS: computer application system; system; treatment;  
storage system; method; digital computer; computer; hardware; computer  
application; utilization; simulation

CLASSIFICATION CODE(S): JD03030U; JC020100

10/5/16 (Item 2 from file: 94)

DIALOG(R)File 94:JICST-EPlus

(c)2004 Japan Science and Tech Corp(JST). All rts. reserv.

02615029 JICST ACCESSION NUMBER: 94A0054587 FILE SEGMENT: JICST-E  
Dynamic Skew Handling for Hash-Joins in a Shared-Nothing Database System.

DAVIS S (1); KITSUREGAWA M (2)

(1) E. Resu Lab.; (2) Univ. Tokyo

Denshi Joho Tsushin Gakkai Shinpojiumu Ronbunshu, 1993, VOL.93,NO.9, PAGE.113-122,  
FIG.1, TBL.4, REF.10

JOURNAL NUMBER: Y0978BAT

UNIVERSAL DECIMAL CLASSIFICATION: 681.3:061.68 681.32

LANGUAGE: English COUNTRY OF PUBLICATION: Japan

DOCUMENT TYPE: Conference Proceeding

ARTICLE TYPE: Original paper

MEDIA TYPE: Printed Publication

ABSTRACT: When data are uniformly distributed, parallel hash-based join  
algorithm scales up well. However, the presence of data skew can cause  
load unbalances among the processors, significantly deteriorating its  
performance. Within the last years, there has been a growing interest  
in addressing the problem of data skew. Many proposed algorithms add an  
extra sampling or scanning phase and a scheduling phase to the usual  
hash and join phases. Based on statistics of fine partitions of the  
relations Obtained in the sampling or scanning phase, they make  
estimations of the join costs in the scheduling phase, and attempt to  
balance the load across the multiple processors for the subsequent join  
phase. However, these algorithms still rely on simplistic assumptions  
of the data distribution within each fine partition. The optimality of  
the resulting load balance depends on the validity of these

assumptions, which, in some cases, results in discrepancies with the real data and thus, in an unexpected join product skew. In this paper we propose a dynamic skew handling algorithm which deals with these load unbalances, by detecting and handling at run-time the join product skew, which could not be correctly predicted in the scheduling phase. When an overloaded processor is detected at run-time, the overload of producing the skewed output tuples (mainly writing the large number of result tuples) is dynamically migrated to other processors, while the reading of the building and probing partitions are stuck at the overloaded processor. We propose two new methods, the result redistribution and the processing task migration, to handle this overload migration. Simulation results show that our dynamic skew handling approach can detect and handle load unbalances efficiently, so that the rebalance of load among the processors results in an almost constant join execution time under different join product skews.  
(author abst.)

DESCRIPTORS: **database** ; hashing; probability distribution; parallel processing; algorithm; multiprocessor system; query processing; performance analysis

BROADER DESCRIPTORS: storage system; method; distribution; treatment; computer system(hardware); system; information processing; analysis

CLASSIFICATION CODE(S): JD030300; JC020100

10/5/17 (Item 1 from file: 144)  
DIALOG(R) File 144:Pascal  
(c) 2004 INIST/CNRS. All rts. reserv.

13946466 PASCAL No.: 99-0129504

**Performance of load balancing techniques for join operations in shared-nothing database management systems**

HUA K A; TAVANAPONG W; LO Y L

School of Computer Science, University of Central Florida, Orlando, Florida 32816-2362, United States; Department of Information Management, Chao Yang University of Technology, TaiChung County, China

Journal: Journal of parallel and distributed computing, 1999, 56 (1)  
17-46

ISSN: 0743-7315 Availability: INIST-20948; 354000074185940020

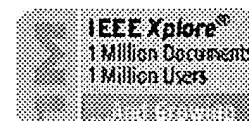
No. of Refs.: 26 ref.

Document Type: P (Serial) ; A (Analytic)

Country of Publication: United States

Language: English

We investigate various load balancing approaches for hash-based join techniques popular in multicomputer-based shared-nothing database systems. When the tuples are not uniformly distributed among the hash buckets, redistribution of these buckets among the processors is necessary to maintain good system performance. Two recent load balancing techniques which rely on sampling and incremental balancing, respectively, have been shown to be more robust than conventional methods. The comparison of these two approaches, however, has not been investigated. In this study, we improve these two schemes and implement them along with a conventional method and a standard join technique which does not do load balancing on an nCUBE/2 parallel computer to compare their performance. Our experimental results indicate that the sampling technique is the better approach. To further evaluate the performance of these techniques under diverse hardware conditions, we also develop a cost model and implement a simulator to perform sensitivity analyses with respect to various hardware parameters. The simulation results show that both sampling and incremental techniques provide noticeable savings over conventional methods, with the sampling approach being more scalable in supporting very large database systems.

[IEEE HOME](#) | [SEARCH IEEE](#) | [SHOP](#) | [WEB ACCOUNT](#) | [CONTACT IEEE](#)[Membership](#) | [Publications/Services](#) | [Standards](#) | [Conferences](#) | [Careers/Jobs](#)**IEEE Xplore**  
RELEASE 1.8Welcome  
United States Patent and Trademark Office[Help](#) | [FAQ](#) | [Terms](#) | [IEEE Peer](#)[Quick Links](#)[» Search Results](#)[Review](#)

Welcome to IEEE Xplore®

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced
- ☐ CrossRef

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

IEEE Enterprise

- ☐ Access the IEEE Enterprise File Cabinet

Print Format

Your search matched **0** of **1128145** documents.  
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

**Refine This Search:**

You may refine your search by editing the current search expression or entering a new one in the text box.

☐ Check to search within this result set**Results Key:****JNL** = Journal or Magazine   **CNF** = Conference   **STD** = Standard**Results:****No documents matched your query.**

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC](#) | [Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

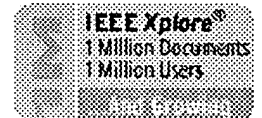
Copyright © 2004 IEEE — All rights reserved

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE



Membership | Publications/Services | Standards | Conferences | Careers/Jobs

**IEEE Xplore**  
 RELEASE 1.8

 Welcome  
 United States Patent and Trademark Office

[Help](#) | [FAQ](#) | [Terms](#) | [IEEE Peer](#) | [Quick Links](#)
» [Advanced Search](#)

Welcome to IEEE Xplore®

- [Home](#)
- [What Can I Access?](#)
- [Log-out](#)

Tables of Contents

- [Journals & Magazines](#)
- [Conference Proceedings](#)
- [Standards](#)

Search

- [By Author](#)
- [Basic](#)
- [Advanced](#)
- [CrossRef](#)

Member Services

- [Join IEEE](#)
- [Establish IEEE Web Account](#)
- [Access the IEEE Member Digital Library](#)

IEEE Enterprise

- [Access the IEEE Enterprise File Cabinet](#)

**Try our New Full-text Search Prototype** **GO**
[Help](#)

- 1) Enter a single keyword, phrase, or Boolean expression.  
Example: acoustic imaging (means the phrase acoustic imaging plus any stem variations)
- 2) Limit your search by using search operators and field codes, if desired.  
Example: optical <and> (fiber <or> fibre) <in> ti
- 3) Limit the results by selecting Search Options.
- 4) Click Search. See [Search Examples](#)

 (hash <sentence> join) <and>  
 (split <sentence> vector)

Start Search

Clear

Note: This function returns plural and suffixed forms of the keyword(s).

 Search operators: <and> <or> <not> <in> [More](#)

 Field codes: au (author), ti (title), ab (abstract), jn (publication name), de (index term) [More](#)
**Search Options:****Select publication types:**

- ☒ IEEE Journals
- ☒ IEE Journals
- ☒ IEEE Conference proceedings
- ☒ IEE Conference proceedings
- ☒ IEEE Standards

**Select years to search:**
 From year:  to 
**Organize search results by:**Sort by: In:  orderList  Results per page
[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC](#) | [Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved



US Patent &amp; Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

+hash +join +split +vector



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **hash join split vector**Found **185** of **150,138**

Sort results by

relevance

[Save results to a Binder](#)[Try an Advanced Search](#)

Display results

expanded form

[Search Tips](#)[Try this search in The ACM Guide](#)
☐ Open results in a new window

Results 1 - 20 of 185

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

# 1 [Research sessions: query processing I: A scalable hash ripple join algorithm](#)

Gang Luo, Curt J. Ellmann, Peter J. Haas, Jeffrey F. Naughton

 June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.12 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recently, Haas and Hellerstein proposed the hash ripple join algorithm in the context of online aggregation. Although the algorithm rapidly gives a good estimate for many join-aggregate problem instances, the convergence can be slow if the number of tuples that satisfy the join predicate is small or if there are many groups in the output. Furthermore, if memory overflows (for example, because the user allows the algorithm to run to completion for an exact answer), the algorithm degenerates to bl ...

# 2 [Query evaluation techniques for large databases](#)

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available: pdf(9.37 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

# 3 [A performance evaluation of four parallel join algorithms in a shared-nothing multiprocessor environment](#)

Donovan A. Schneider, David J. DeWitt

 June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: pdf(1.48 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


In this paper we analyze and compare four parallel join algorithms. Grace and Hybrid hash

represent the class of hash-based join methods, Simple hash represents a looping algorithm with hashing, and our last algorithm is the more traditional sort-merge. The performance of each of the algorithms with different tuple distribution policies, the addition of bit vector filters, varying amounts of main-memory for joining, and non-uniformly distributed join attribute values is studied. The Hybrid ...

4 GESS: a scalable similarity-join algorithm for mining large data sets in high dimensional spaces

Jens-Peter Dittrich, Bernhard Seeger

August 2001 **Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining**


Full text available:  pdf(794.22 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The similarity join is an important operation for mining high-dimensional feature spaces. Given two data sets, the similarity join computes all tuples  $(x, y)$  that are within a distance  $\epsilon$ . One of the most efficient algorithms for processing similarity-joins is the Multidimensional-Spatial Join (MSJ) by Koudas and Sevcik. In our previous work --- pursued for the two-dimensional case --- we found however that MSJ has several performance shortcomings in terms of CPU and I/O cost as ...

5 Parallel main memory database system

Soon M. Chung


April 1992 **Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing: technological challenges of the 1990's**

Full text available:  pdf(946.35 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

6 A performance analysis of the gamma database machine

D. J. DeWitt, S. Ghanderaizadeh, D. Schneider

June 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international conference on Management of data**, Volume 17 Issue 3

Full text available:  pdf(1.42 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the results of an initial performance evaluation of the Gamma database machine. In our experiments we measured the effect of relation size and indices on response time for selection, join, and aggregation queries, and single-tuple updates. A Teradata DBC/1012 database machine of similar size is used as a basis for interpreting the results obtained. We also analyze the performance of Gemma relative to the number of processors employed and study the impact of varying the m ...

7 Adaptive algorithms for set containment joins

Sergey Melnik, Hector Garcia-Molina

March 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 1

Full text available:  pdf(465.76 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A set containment join is a join between set-valued attributes of two relations, whose join condition is specified using the subset ( $\subseteq$ ) operator. Set containment joins are deployed in many database applications, even those that do not support set-valued attributes. In this article, we propose two novel partitioning algorithms, called the Adaptive Pick-and-Sweep Join (APSJ) and the Adaptive Divide-and-Conquer Join (ADCJ), which allow computing set containment joins efficiently. We show that ...

8 Functional-join processing

R. Braumandl, J. Claussen, A. Kemper, D. Kossmann

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4.

Full text available:  [pdf\(486.22 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)


Inter-object references are one of the key concepts of object-relational and object-oriented database systems. In this work, we investigate alternative techniques to implement inter-object references and make the best use of them in query processing, i.e., in evaluating functional joins. We will give a comprehensive overview and performance evaluation of all known techniques for simple (single-valued) as well as multi-valued functional joins. Furthermore, we will describe special *order-preser ...*

**Keywords:** *Functional join, Logical OID, Object identifier, Order-preserving join, Physical OID, Pointer join, Query processing*

## 9 [Classification: SQL database primitives for decision tree classifiers](#)

Kai-Uwe Sattler, Oliver Dunemann

October 2001 **Proceedings of the tenth international conference on Information and knowledge management**

Full text available:  [pdf\(1.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Scalable data mining in large databases is one of today's challenges to database technologies. Thus, substantial effort is dedicated to a tight coupling of database and data mining systems leading to database primitives supporting data mining tasks. In order to support a wide range of tasks and to be of general usage these primitives should be rather building blocks than implementations of specific algorithms. In this paper, we describe primitives for building and applying decision tree classifi ...

**Keywords:** SQL-aware data mining, data mining primitives, query operators

## 10 [Join processing in relational databases](#)

Priti Mishra, Margaret H. Eich

March 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 1

Full text available:  [pdf\(4.42 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most difficult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples fro ...

**Keywords:** database machines, distributed processing, join, parallel processing, relational algebra

## 11 [High performance data mining \(tutorial PM-3\)](#)

Vipin Kumar, Mohammed Zaki

August 2000 **Tutorial notes of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining**


Full text available:  [pdf\(8.06 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

## 12 [Searching in high-dimensional spaces: Index structures for improving the performance of multimedia databases](#)

Christian Böhm, Stefan Berchtold, Daniel A. Keim

September 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 3

Additional Information:


Full text available:  pdf(1.39 MB)[full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research issue in the field of multimedia databases is the content-based retrieval of similar multimedia objects such as images, text, and videos. However, in contrast to searching data in a relational database, a content-based retrieval requires the search of similar objects as a basic functionality of the database system ...

**Keywords:** Index structures, indexing high-dimensional data, multimedia databases, similarity search

### 13 [Multi-dimensional resource scheduling for parallel queries](#)


Minos N. Garofalakis, Yannis E. Ioannidis

June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data**, Volume 25 Issue 2Full text available:  pdf(1.47 MB)[Additional Information: full citation, abstract, references, citations, index terms](#)

Scheduling query execution plans is an important component of query optimization in parallel database systems. The problem is particularly complex in a shared-nothing execution environment, where each system node represents a collection of time-shareable resources (e.g., CPU(s), disk(s), etc.) and communicates with other nodes only by message-passing. Significant research effort has concentrated on only a subset of the various forms of intra-query parallelism so that scheduling and synchronizati ...

### 14 [Multidimensional access methods](#)

Volker Gaede, Oliver Günther


June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2Full text available:  pdf(1.05 MB)[Additional Information: full citation, abstract, references, citations, index terms](#)

Search operations in databases require special support at the physical level. This is true for conventional databases as well as spatial databases, where typical search operations include the point query (find all objects that contain a given search point) and the region query (find all objects that overlap a given search region). More than ten years of spatial database research have resulted in a great variety of multidimensional access methods to support ...

**Keywords:** data structures, multidimensional access methods

### 15 [Parallel execution of prolog programs: a survey](#)

Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo

July 2001 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 23 Issue 4Full text available:  pdf(1.95 MB)[Additional Information: full citation, abstract, references, citations, index terms](#)

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, the presence of nondeterminism, and their referential transparency, among other characteristics, make logic programs interesting candidates for obtaining speedups through parallel execution. At the same time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

**Keywords:** Automatic parallelization, constraint programming, logic programming, parallelism, prolog

## 16 XML indexing and compression: Efficient processing of joins on set-valued attributes

Nikos Mamoulis

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(678.13 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Object-oriented and object-relational DBMS support set valued attributes, which are a natural and concise way to model complex information. However, there has been limited research to-date on the evaluation of query operators that apply on sets. In this paper we study the join of two relations on their set-valued attributes. Various join types are considered, namely the set containment, set equality, and set overlap joins. We show that the inverted file, a powerful index for selection queries, c ...

## 17 External memory algorithms and data structures: dealing with

# massive data

Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Full text available:  pdf(628.46 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

**Keywords:** B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

## 18 Object-based and image-based object representations

Hanan Samet

June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2

Full text available:  pdf(1.05 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

**Keywords:** Access methods, R-trees, feature query, geographic information systems (GIS), image space, location query, object space, octrees, pyramids, quadrees, space-filling curves, spatial databases

## 19 Technique for automatically correcting words in text

Karen Kukich

December 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 4

Full text available:  pdf(6.23 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)


Research aimed at correcting words in text has focused on three progressively more difficult problems: (1) nonword error detection; (2) isolated-word error correction; and (3) context-dependent word correction. In response to the first problem, efficient pattern-matching and n-gram analysis techniques have been developed for detecting strings that do not appear in a given word list. In response to the second problem, a variety of general and application-specific spelling cor ...

**Keywords:** n-gram analysis, Optical Character Recognition (OCR), context-dependent spelling correction, grammar checking, natural-language-processing models, neural net classifiers, spell checking, spelling error detection, spelling error patterns, statistical-language models, word recognition and correction

## 20 Transformation-based spatial join

Ju-Won Song, Kyu-Young Whang, Young-Koo Lee, Min-Jae Lee, Sang-Wook Kim

November 1999 **Proceedings of the eighth international conference on Information and knowledge management**

Full text available:  pdf (1.51 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)





Spatial join finds pairs of spatial objects having a specific spatial relationship in spatial database systems. A number of spatial join algorithms have recently been proposed in the literature. Most of them, however, perform the join in the original space. Joining in the original space has a drawback of dealing with sizes of objects and thus has difficulty in developing a formal algorithm that does not rely on heuristics. In this paper, we propose a spatial join algorithm based on the tran ...

Results 1 - 20 of 185

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright ?2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.